




FORTINET
High Performance Network Security



FortiOS™ Handbook - Firewall

VERSION 5.4.6



FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

FORTICAST

<http://forticast.fortinet.com>

END USER LICENSE AGREEMENT

<http://www.fortinet.com/doc/legal/EULA.pdf>

FORTINET PRIVACY POLICY

<https://www.fortinet.com/corporate/about-us/privacy.html>

FEEDBACK

Email: techdocs@fortinet.com



Tuesday, January 23, 2018

FortiOS™ Handbook - Firewall

01-540-1248222-20140707

TABLE OF CONTENTS

Change Log	9
Introduction	10
How this Guide is Organized	10
FortiGate Firewall Components	11
How does a FortiGate Protect Your Network?	12
GUI & CLI - What You May Not Know	14
Changing the default column setting on the policy page	14
Example	15
Naming Rules and Restrictions	15
Character Restrictions	16
Object names	16
Passwords	16
Numeric Values	17
To Enable or Disable Optionally Displayed Features	17
What's new for the Firewall in 5.4	19
Learning mode for Firewall policies (310544 365727)	19
New Features in 5.4.1	21
Multiple interfaces or ANY interface can be added to a firewall policy (288984)	21
Multicast policy page changes (293709 305114)	21
Policy objects dialogs updated to new GUI style (354505)	22
New Features in 5.4.0	22
Display change in Policy listing (284027)	22
RPC over HTTP traffic separate (288526)	22
Disable Server Response Inspection supported (274458)	23
Policy counter improvements (277555 260743 172125)	23
Bidirectional Forwarding Detection (BFD) (247622)	23
TCP sessions can be created without TCP syn flag checking (236078)	23
Mirroring of traffic decrypted by SSL inspection (275458)	24
Support for full cone NAT (269939)	24
Enable or disable inspecting IPv4 and IPv6 ICMP traffic (258734)	24
Policy names (246575 269948 293048)	25
Policy and route lookup (266996 222827)	25
Support NAT 64 CLAT (244986)	26
VIPs can contain FQDNs (268876)	26

Access Control Lists (ACLs) (293399).....	26
GUI improvement for DoS Policy configuration (286905).....	27
Expired Policy Object warnings (259338).....	27
Firewall concepts.....	28
What is a Firewall?.....	28
Network Layer or Packet Filter Firewalls.....	28
Application Layer Firewalls.....	29
Proxy Servers.....	30
Security Profiles.....	30
FortiGate Modes.....	31
NAT/Route Mode.....	31
Transparent Mode.....	32
How Packets are handled by FortiOS.....	32
Interfaces and Zones.....	33
Access Control Lists.....	34
Incoming Interfaces.....	34
Addresses.....	34
Services.....	35
IPv6.....	35
IPv6 in FortiOS.....	36
Dual Stack routing configuration.....	37
IPv6 Tunneling.....	37
Tunneling IPv6 through IPsec VPN.....	38
NAT.....	38
The Origins of NAT.....	39
Dynamic NAT.....	39
Static NAT.....	40
Benefits of NAT.....	41
NAT in Transparent Mode.....	42
Central NAT Table.....	43
NAT 64 and NAT46.....	43
NAT 66.....	44
How FortiOS differentiates sessions when NATing.....	45
IP Pools.....	52
Source IP address and IP pool address matching when using a range.....	53
Scenario 1:.....	53
Scenario 2:.....	53
Scenario 3:.....	53
ARP Replies.....	53
IP pools and zones.....	54
Fixed Port.....	54
Match-VIP.....	54

Services and TCP ports	54
Protocol Types	55
Protocol Port Values	58
ICMP	58
ICMPv6	63
IP	67
Security policies	82
Firewall policies	82
Firewall policy parameters	83
What is not expressly allowed is denied	85
Policy order	85
Policy Identification	88
UUID Support	88
Security profiles	88
AntiVirus	89
Web Filtering	89
Application Control	89
Intrusion Protection (IPS)	89
Email Filtering	90
Data Leak Prevention (DLP)	90
VoIP	90
ICAP	90
EndPoint Control	90
Proxy Option Components	91
The use of different proxy profiles and profile options	91
Viewing Firewall Policies	93
How “Any” policy can remove the Section View	94
Security policy configuration extensions	95
SSL/SSH Inspection	95
Inspection Exemption	96
Identity Based Policies	98
Identity-based policy positioning	98
VPN Policies	99
IPsec Policies	100
DoS Protection	100
One-Arm IDS	104
IPv6 IPS	104
Traffic Destined to the FortiGate unit	104
Dropped, Flooded, Broadcast, Multicast and L2 packets	104
GUI and CLI	105
Local-In Policies	105

Security Policy 0.....	106
Deny Policies.....	106
Accept Policies.....	107
Fixed Port.....	107
Endpoint Security.....	107
Traffic Logging.....	108
Quality of Service.....	109
Queuing.....	109
Policy Monitor.....	110
Upper Pane.....	110
Lower Pane.....	110
Network defense.....	112
Monitoring.....	112
Blocking external probes.....	112
Address sweeps.....	112
Port scans.....	113
Probes using IP traffic options.....	113
Evasion techniques.....	114
Defending against DoS attacks.....	117
The “three-way handshake”.....	117
SYN flood.....	117
SYN spoofing.....	118
DDoS SYN flood.....	119
Configuring the SYN threshold to prevent SYN floods.....	120
SYN proxy.....	120
Other flood types.....	120
DoS policies.....	120
Inside FortiOS: Denial of Service (DoS) Protection.....	122
About DoS and DDoS attacks.....	122
FortiOS DoS and DDoS protection.....	122
FortiOS DDoS Prevention.....	123
Configuration options.....	124
Standard configuration.....	124
Out of band configuration (sniffer mode).....	124
DoS policies.....	125
Hardware acceleration.....	125
The FortiGuard Center.....	125
Firewall Policies.....	127
IPv4 Policy.....	127
To configure a IPv4 policy in the GUI.....	127
IPv6 Policy.....	130
To configure a IPv6 policy in the GUI.....	130

NAT64 Policy.....	132
To configure a NAT64 policy in the GUI.....	132
NAT46 Policy.....	134
To configure a NAT46 policy in the GUI.....	134
Central SNAT.....	135
To configure a Central SNAT entry in the GUI.....	136
To configure Central SNAT in the CLI.....	136
Example: Central NAT Table.....	137
IPv4 Access Control List.....	138
To configure a IPv4 Access Control List entry in the GUI.....	138
To configure a IPv4 Access Control List entry in the CLI.....	139
IPv6 Access Control List.....	139
To configure a IPv6 Access Control List entry in the GUI.....	139
To configure a IPv6 Access Control List entry in the CLI.....	140
IPv4 DoS Policy.....	140
To configure a IPv4 DoS Policy in the GUI.....	140
Example.....	141
IPv6 DoS Policy.....	145
To configure a IPv6 DoS Policy in the GUI.....	145
Firewall objects.....	147
Addresses.....	149
Interfaces.....	150
IPv4 Addresses.....	151
FQDN Addresses.....	151
Verification.....	153
Geography Based Addresses.....	153
IP Range Addresses.....	156
IP / Netmask Addresses.....	157
Wildcard FQDN.....	159
IPv6 Addresses.....	160
Subnet Addresses.....	160
Multicast Addresses.....	161
Multicast IP Range.....	161
Broadcast Subnet.....	163
Multicast IP addresses.....	163
Explicit Proxy Addresses.....	164
Address Groups.....	167
UUID Support.....	168
Virtual IPs.....	169
Creating a Virtual IP.....	170
Dynamic VIP according to DNS translation.....	172
Virtual IP Groups.....	173

Creating a Virtual IP Group.....	173
Configuring IP pools.....	173
Creating a IPv4 Pool.....	174
Creating a IPv6 Pool.....	177
Services.....	178
Categories.....	179
Configuring a new service.....	180
Example Scenario: Using FortiGate services to support Audio/Visual Conferencing....	184
Specific Addresses in TCP/UDP/SCTP.....	191
Service Groups.....	191
Firewall schedules.....	193
Creating a recurring schedule object.....	193
Creating a One-time schedule object.....	194
Schedule expiration.....	195
Firewall-session-dirty setting.....	195
Schedule Groups.....	196
Creating a recurring schedule object.....	196
Schedule Expiration.....	196

Change Log

Date	Change Description
2017-12-07	Correction made to Proxy options components in Security profiles section.
2017-11-01	Additional content added to geography addresses and packet defragmentation
2017-09-21	Caution added to recurring schedule objects.
2017-05-3	Additional content added to DoS Protection
2017-02-21	Changes to Wildcard FQDN
2017-02-06	Character restrictions edit
2016-12-19	Static Route Configuration
2014-07-07	Initial release.

Introduction

Welcome and thank you for selecting Fortinet products for your network protection. This document is intended to provide the concepts and techniques that will be needed to configure the FortiGate firewall on your FortiGate unit.

Before you start administrating your FortiGate device, certain assumptions have been made in the writing of this manual:

- You have administrative access to the Web based GUI or to the Command Line Interface.
- The FortiGate unit is integrated into your network.
- The operation mode (NAT or Transparent) has been configured.
- Network Interfaces have been configured.
- DNS settings have been configured.
- The system time settings have been configured.
- Firmware is up to date.
- FortiGuard Service licences are current and the device is able to connect to the FortiGuard Servers.
- If you are using FortiCloud, it is properly configured.

How this Guide is Organized

"Firewall concepts" explains the ideas behind the components, techniques and processes that are involved in setting up and running a firewall in general and the FortiGate firewall in particular. The premise here is that regardless of how experienced someone is with firewalls as they go through the process of configuring a firewall that is new to them they are likely to come across a term or setting that they may not be familiar with even if it is only in the context of the setting they are working in at the moment. FortiGate firewall are quite comprehensive and can be very granular in the functions that they perform, so it makes sense to have a consistent frame of reference for the ideas that we will be working with.

Some examples of the concepts that will be addressed here are:

- "What is a Firewall?"
- "NAT"
- "IPv6"

"Firewall objects" describes the following firewall objects:

- Addressing
- Services
- Firewall Policies

"Network defense" describes various methods of defending your Network using the abilities of the FortiGate Firewall.

"GUI & CLI - What You May Not Know" helps you navigate and find the components in the Web-based Manager that you will need to build the functions. This section does not include any in-depth explanations of what each

object does as that is covered in the concepts section. This section is for showing you where you need to input your information and let you know what format the interface expects to get that information

"Building firewall objects and policies" is similar to a cookbook in that it will refer to a number of common tasks that you will likely perform to get the full functionality out of your FortiGate firewall. Because of the way that firewall are designed, performing many of the tasks requires that firewall components be set up in a number of different sections of the interface and be configured to work together to achieve the desired result. This section will bring those components all together as a straight forward series of instructions.

"Multicast forwarding" is a reference guide including the concepts and examples that are involved in the use of multicast addressing and policy forwarding as it is used in the FortiGate firewall.

FortiGate Firewall Components

The FortiGate firewall is made up of a number of different components that are used to build an impressive list of features that have flexibility of scope and granularity of control that provide protection that is beyond that provided by the basic firewalls of the past.

Some of the components that FortiOS uses to build features are:

- Interfaces
- VLANs
- Soft Switches
- Zones
- Predefined Addresses
- IP address based
- FQDN based
- Geography based
- Access Schedules
- Authentication
- Local User based
- Authentication Server based (Active Directory, Radius, LDAP)
- Device Based
- Configureable Services
- IPv4 and IPv6 protocol support

The features of FortiOS include but are not limited to:

- Security profiles, sometimes referred to as Unified Threat Management (UTM) or Next Generation Firewall (NGFW)
- Predefined firewall addresses (this includes IPv4 and IPv6, IP pools, . wildcard addresses and netmasks, and geography-based addresses)
- Monitoring traffic
- Traffic shaping and per-IP traffic shaping (advanced)
- Firewall schedules
- Services (such as AOL, DHCP and FTP)
- Logging traffic
- Quality of Service (QoS)

- Identity-based policies
- Endpoint security

The "Firewall concepts" expand on what each of the features does and how they relate to the administration of the FortiGate firewall. The section will also try to explain some of the common firewall concepts that will be touched on in the implementing of these features.

"Building firewall objects and policies" shows how to perform specific tasks with the FortiGate firewall.

How does a FortiGate Protect Your Network?

The FortiGate firewall protects your network by taking the various components and using them together to build a kind of wall or access control point so that anyone that is not supposed to be on your network is prevented from accessing your network in anyway other than those approved by you. It also protects your network from itself by keeping things that shouldn't happen from happening and optimizing the flow of traffic so that the network is protected from traffic congestion that would otherwise impede traffic flow.

Most people have at one time or another played with the children's toy system that is made up of interlocking blocks. The blocks come in different shapes and sizes so that you can build structures to suit your needs and in your way. The components of the FortiGate firewall are similar. You are not forced to use all of the blocks all of the time. You mix and match them to get the results that you are looking for. You can build a very basic structure that's only function is to direct traffic in and out to the correct subnets or you can build a fortress that only allows specific traffic to specific hosts from specific hosts at specific times of day and that is only if they provide the credentials that have been pre-approved and all of the traffic is encrypted so that even when the traffic is out on the Internet it is private from the world. Just like the interlocking blocks, what you build is up to you, but chances are if you put them together the right way there isn't much that can't be built.

Here is one example of how the components could be put together to support the requirements of a network infrastructure design.

- Off the Internal interface you could have separate VLANs. One for each for the departments of Sales, Marketing and Engineering so that the traffic from the users on one VLAN does not intrude upon the hosts of the other VLANs and the department are isolated from one another for security reasons.
- To ease in the administration each of the VLAN sub-interfaces is made a member of a zone so that security policies that apply to all of the hosts on all of the VLANs can be applied to all of them at once.
- Using the addresses component each of the IP address ranges could be assigned a user friendly name so that they could be referred to individually and then for policies that would refer to them all as a whole the individual ranges to be made members of an address group.
- Firewall schedules could be created to address the differing needs of each of the groups so that Sales and Marketing could be allowed access to the Internet during regular business hours and the Engineering department could be allowed access during the lunch break.
- By setting up the outgoing policies to use FortiGuard Web-filtering the employees could be prevented from visiting inappropriate sites and thus enforcing the policies of the HR department.
- A couple of virtual IP addresses with port forwarding could be configured to allow users on the Internet to access a web server on the DMZ subnet using the company's only Public IP address without affecting the traffic that goes to the company's mail server that is hosted on a complete different computer.
- Even though the Web server on the same DMZ has an FTP service to allow for the uploading of web pages to the web server from the Marketing and Engineer teams, by placing a DENY policy on any FTP traffic from the Internet malicious users are prevented from abusing the FTP service.

- By monitoring the traffic as it goes through the policies you can verify that the policies are in working order.
- By using a combination of ALLOW and DENY policies and placing them in the correct order you could arrange for an outside contractor to be allowed to update the web site as well

These set of configurations is not extensive but it does give an idea of how different components can be mixed and matched to build a configuration that meets an organization's needs but at the same time protect it from security risks.

GUI & CLI - What You May Not Know

The Graphic User Interface (GUI) is designed to be as intuitive as possible but there are always a few things that are left out because to put all of that information on the interface would clutter it up to the point where it wouldn't be graphical and intuitive anymore.

This section is made up of knowledge that will make working with the both of the management interfaces easier because you won't have to find out about things like field limitations through trial and error. Some of it has to do with changing in how navigation in the GUI has changed.

The section includes the topics:

- Mouse Tricks
- Changing the default column setting on the policy page
- Naming Rules and Restrictions
- Character Restrictions
- Length of Fields Restrictions
- Object Tagging and Coloring
- Numeric Values
- Selecting options from a list
- Enabling or disabling options
- To Enable or Disable Optionally Displayed Features

Changing the default column setting on the policy page

The **Policy & Objects > Policy > IPv4** window is one of the more important ones in the Web based interface and has the capacity to display a lot of information, but displaying all of that information at the same time makes for a very busy screen. If all of the columns are displayed, depending on the screen size you may have to constantly use the scroll bars to see what you need to look at. The default installation shows some of the more commonly used columns but these list may not consist of the columns that you wish to look at or the order that you wish to view them in. For this reason it is possible, through the CLI to override these settings to establish a new default.

The syntax of the command starts with:

```
config system settings
    set gui-default-policy-columns
```

The rest of the command is a space delimited list that depends on the columns you wish to view and the order you wish to view them in. The possible selection is in the following table.

Variables for the `gui-default-policy-columns` command

Variable Name	Column Heading
#	Sequence Number
policyid	Policy ID
srcintf	Source Interface
dstintf	Destination Interface
srcaddr	Source Addresses
dstaddr	Destination Addresses
schedule	Policy Schedule
service	Policy Services
action	Policy Action
logtraffic	Traffic Logging Status
nat	Policy NAT Status
status	Policy Status
authentication	Authentication Groups
count	Policy Traffic Counter
profile	Security Profiles
vpntunnel	VPN Tunnel
comments	Policy Comment

Example

If you wanted these columns in this order, Policy ID, Source Addresses, Destination Addresses, Security Profiles, Policy Comment. You would enter the command:

```
config system settings
set gui-default-policy-columns policyid srcaddr dstaddr profile comments
```

Naming Rules and Restrictions

The following are the specific rules that are obeyed by the FortiGate.

Duplicate Name Issues:

- A VLAN cannot have the same name as a physical interface.
- An Address must not have the same name as an Address Group.
- An Address or Address Group must not have the same name as a Virtual IP Address.
- A Service cannot have the same name as a Service Group.
- A VLAN must not have the same name as a VDOM.
- A VLAN or VDOM must not have the same name as a Zone.



Try to make each firewall object name as unique as possible so that it cannot be confused with another object.

Character Restrictions

Object names

Object names cannot contain the following characters:

- <
- >
- (
-)

All other type-able characters, including other languages, are supported as usable characters.



FortiOS allows spaces in just about all object name fields, but this was not always the case. Some people are still cautious where this is concerned. If you're cautious about characters in the names of objects use this basic rule of thumb:

When naming objects, only use characters that are alphanumeric (a-z, A-Z, 0-9) and where there is the temptation to use spaces in a name, use the '-' (dash) and '_' (underscore).

Passwords

Password fields don't have the same character restrictions as the object names, but there is one restriction. Only type-able characters can be used as a password.



The idea of non type-able characters may sound strange, but in some systems it is theoretically possible, using ASCII codes to input icons or even null characters into a field. These are not accepted in FortiOS.

Numeric Values

Numeric values are used to configure various sizes, rates, numeric addresses, or other numeric values. For example, a static routing priority of 10, a port number of 8080, or an IP address of 10.10.10.1. Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or as in the case of MAC or IPv6 addresses separated by colons (for example, the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base-10 numbers, but some fields (again such as MAC addresses) require hexadecimal numbers.

Most web-based manager numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help includes information about allowed numeric value ranges. Both the web-based manager and the CLI prevent you from entering invalid numbers.

To Enable or Disable Optionally Displayed Features

There are a number of features in the web-based manager that can be configured to either be displayed if you are likely to use them or disabled if you have no need to see them. Some of the ones that may be relevant to the function of the Firewall are:

- Central NAT Table
- Dynamic Profile
- Explicit Proxy
- Implicit Firewall Policies
- IPv6
- Local In Policy
- Multicast Policies
- Allow unnamed Policies
- Allow Multiple Interfaces in policies

You can enable or disable these features by going to **System > Admin > Settings**. You can also change this settings in the CLI. A number of options for this can be found in system settings or by using the following CLI options:

```
config system settings
  set central-nat [disable | enable]
  set gui-icap [disable | enable]
  set gui-nat46-64 [disable | enable]
  set gui-dns-database [disable | enable]
  set gui-load-balance [disable | enable]
  set gui-multicast-policy [disable | enable]
  set gui-local-in-policy [disable | enable]
  set gui-local-reports [disable | enable]
  set gui-explicit-proxy [disable | enable]
  set gui-policy-based-ipsec [disable | enable]
  set gui-ips [disable | enable]
  set gui-endpoint-on-net [disable | enable]
  set gui-waf-profile [disable | enable]
  set gui-allow-unnamed-policy [disable | enable]
  set gui-multiple-interface-policy [disable | enable]
```

end

What's new for the Firewall in 5.4

Learning mode for Firewall policies (310544 365727)

The learning mode feature is a quick and easy method for setting a policy to allow everything but to log it all so that it can later be used to determine what restrictions and protections should be applied. The objective is to monitor the traffic not act upon it while in Learning mode.

Once the **Learn** action is enabled, functions produce hard coded profiles that will be enabled on the policy. The following profiles are set up:

- AntiVirus (av-profile)
- Web Filter (webfilter-profile)
- Anti Spam(spamfilter-profile)
- Data Leak Prevention (dlp-sensor)
- Intrusion Protection (ips-sensor)
- Application Control (application-list)
- Proxy Options (profile-protocol-options)



- These UTM profiles are all using Flow mode
 - SSL inspection is always disable for the Learn option
 - These profiles are static and cannot be edited.
-

Profiles that are not being used are:

- DNS Filter (Does not have a Flow mode)
- Web Application Firewall(Does not have a Flow mode)
- CASI(Almost all signatures in CASI require SSL deep inspection. Without SSL inspection, turning on CASI serves little purpose)

The ability to allow policies to be set to a learning mode is enabled on a per VDOM basis.

```
config system settings
  set gui-policy-learning [enable | disable]
end
```

Once the feature is enabled on the VDOM, Learn is an available **Action** option when editing a policy.



Because this feature requires a minimum level of logging capabilities, it is only available on FortiGates with hard drives. Smaller models may not be able to use this feature.

New Policy

Name

Incoming Interface

+

Outgoing Interface

+

Source

+

Destination Address

+

Schedule

always

Service

+

Action

✓ ACCEPT

✗ DENY

LEARN

IPsec

Firewall / Network Options

NAT

☒

Comments

Write a comment...

0/1023

Enable this policy

☒

OK

Cancel

Once the Learning policy has been running for a sufficient time to collect needed information a report can be looked at by going to **Log & Report > Learning Report**.

The Report can be either a **Full Report** or a **Report Summary**

The time frame of the report can be **5 minutes**, **1 hour**, or **24 hours**.

The Learning Report includes:

Deployment Methodology

- Test Details
 - Start time
 - End time
 - Model
 - Firmware
- Policy List

Executive Summary

- Total Attacks Detected
- Top Application Category
- Top Web Category
- Top Web Domain
- Top Host by Bandwidth
- Host with Highest Session Count

Security and Threat Prevention

- High Risk Applications
- Application Vulnerability Exploits

- Malware, botnets and Spyware/Adware
- At-Risk Devices and Hosts

User Productivity

- Application Usage
 - Top Application Categories
 - Top Social Media Applications
 - Top Video/Audio Streaming Applications
 - Top Peer to Peer Applications
 - Top Gaming Applications
- Web Usage
 - Top Web Categories
 - Top Web Applications
 - Top Web Domains

New Features in 5.4.1

Multiple interfaces or ANY interface can be added to a firewall policy (288984)

This feature can be enabled or disabled in the GUI by going to the **System > Feature Select** page and toggling **Multiple Interface Policies**.

When selecting the **Incoming** or **Outgoing** interface of a policy, there are a few choices:

- The ANY interface (choosing this will remove all other interfaces)
- A single specific interface
- multiple specific interfaces (can be added at the same time or one at a time)

The GUI is intuitive and straightforward on how to do this. Click on the "+" symbol in the interface field and then select the desired interfaces from the side menu. There are a couple of ways to do it in the CLI:

1. Set the interfaces all at once:

```
config firewall policy
edit 0
set srcintf wan1 wan2
end
```





2. Set the first interface and append additional ones:

```
config firewall policy
edit 0
set srcintf wan1
append srcintf wan2
end
```

Multicast policy page changes (293709 305114)

The multicast policy GUI page has been updated to the new GUI look and feel. Some functionality has also been changed.

- The DNAT option has been removed from the GUI but is still in the CLI, you can set the action to IPsec, and if you select Log Allowed Traffic you can also select a few logging options.
- The Multicast policy page loads faster.

Incoming Interface	 port1	▼
Outgoing Interface	 port2	▼
Source Address	 all	✕
Destination Address	 all	✕
Action	<div>ACCEPT</div> <div>DENY</div> <div>IPsec</div>	
Enable SNAT	<input type="checkbox"/>	
Protocol	Any ▼	
<input checked="" type="checkbox"/> Log Allowed Traffic		
Generate Logs when Session Starts <input checked="" type="checkbox"/>		
Capture Packets <input type="checkbox"/>		
Enable this policy <input checked="" type="checkbox"/>		

Policy objects dialogs updated to new GUI style (354505)

To avoid confusion, the default value for "day" is no longer Sunday. In the GUI, none of the day options are selected.

New Features in 5.4.0

Display change in Policy listing (284027)

Alias names for interfaces, if used now appear in the headings for the Interface Pair View or what used to be called the Section View.

RPC over HTTP traffic separate (288526)

How protocol options profiles and SSL inspection profiles handle RPC (Remote Procedure Calls) over HTTP traffic can now be configured separately from normal HTTP traffic.

CLI syntax changes

```
config firewall profile-protocol-options
  edit 0
    set rpc-over-http {disable | enable}
  end

config firewall ssl-ssh-profile
  edit deep-inspection
    set rpc-over-http {disable | enable}
  end
```

Disable Server Response Inspection supported (274458)

Disable Server Response Inspection (DSRI) option included in Firewall Policy (CLI only) to assist performance when only using URL filtering as it allows the system to ignore the http server responses.

CLI syntax for changing the status of the DSRI setting:

```
conf firewall policy|policy6
  edit NNN
    set dsri enable/disable
  end

conf firewall interface-policy|interface-policy6
  edit NNN
    set dsri enable/disable
  end

conf firewall sniffer
  edit NNN
    set dsri enable/disable
  end
```

Policy counter improvements (277555 260743 172125)

- implicit deny policy counter added
- first-hit time tracked for each policy
- "Hit count" is tracked for each policy (total number of new sessions since last reset)
- Most counters now persist across reboots

Bidirectional Forwarding Detection (BFD) (247622)

Bidirectional Forwarding Detection (BFD) protocol support has been added to Protocol Independent Multicast (PIM), to detect failures between forwarding engines.

TCP sessions can be created without TCP syn flag checking (236078)

A Per-VDOM option is available to enable or disable the creation of TCP sessions without TCP SYN flag checking

Mirroring of traffic decrypted by SSL inspection (275458)

This feature sends a copy of traffic decrypted by SSL inspection to one or more FortiGate interfaces so that it can be collected by raw packet capture tool for archiving and analysis.

This feature is available if the inspection mode is set to flow-based. Use the following command to enable this feature in a policy. The following command sends all traffic decrypted by the policy to the FortiGate port1 and port2 interfaces.

```
conf firewall policy
edit 1
set ssl-mirror enable/disable
set ssl-mirror-intf port1 port2
next
```

Support for full cone NAT (269939)

Full cone NAT maps a public IP address and port to a LAN IP address and port. This means that a device on the Internet can send data to the internal LAN IP address and port number by directing it to the external IP address and port number. Sending to the correct IP address but a different port will cause the communication to fail. This type of NAT is also known as port forwarding.

Full cone NATing is configured only in the CLI. It is done by properly configuring an IP pool for the NATing of an external IP address. The two important settings are:

- `set type` - it must be set to `port-block-allocation` to use full cone
- `set permit-any-host` - enabling it is what enables full cone NAT

An example for the IP pool configuration would be:

```
config firewall ippool
edit "full_cone-pool1"
set type port-block-allocation
set startip 10.1.1.1
set endip 10.1.1.1
set permit-any-host enable
end
```

Enable or disable inspecting IPv4 and IPv6 ICMP traffic (258734)

There is now a system setting that determines if ICMP traffic can pass through a Fortigate even if there is no existing session.

```
config system settings
set asymroute-icmp enable
set asymroute6-icmp enable
end
```

When feature enabled:

- Allows ICMP or ICMPv6 reply traffic can pass through firewall when there is no session existing - asymmetric routing case.

When feature disabled:

- Prevents ICMP or ICMPv6 replies from passing through firewall when there is no session existing.

Policy names (246575 269948 293048)

In addition to the Policy ID #, there is now a Policy name field in the policy settings. On upgrading to 5.4, policy names will not be assigned to old policies but when configuring new policies, a unique name must be assigned to it. Every policy name must be unique for the current VDOM regardless of policy type.

In the GUI, the field for the policy name is the first field on the editing page.

In the CLI, the syntax for assigning the policy name is:

```
config firewall [policy|policy6]
    set name <policy_name>
end
```

The feature can be turned on or off.

To turn it off in the CLI:

```
config system settings
    set gui-advance-policy[enable|disable]
end
```

To turn it off in the GUI, the ability to enable or disable it in the GUI must be enabled in the CLI. It is disabled by default. The syntax is:

```
config system settings
    set gui-allow-unnamed-policy [enable | disable]
end
```

Once it has been enabled, the requirement for named passwords can be relaxed by going to **System > Feature Select. Allow Unamed Policies** can be found under **Additional Features**.

This setting is VDOM based so if you are running VDOMs you will have to enter the correct VDOM before entering the CLI commands or turning the feature on or off in the GUI.

Policy and route lookup (266996 222827)

The **Policy Lookup** button in the menu bar at the top of the IPv4 and IPv6 Policy pages is used to determine the policy that traffic with a particular set of parameters will use. Once the parameters are entered, the policy that the traffic will use is displayed.

The parameters are:

- Source Interface - select from drop down menu of available interfaces
- Protocol - select from a drop down menu of:
 - IP
 - TCP
 - UDP
 - SCTP
 - [ICMP|ICMPv6]
 - [ICMP|ICMPv6] ping request
 - [ICMP|ICMPv6] ping reply
- Source - Source IP address
- Source Port
- Destination - Destination IP address

- Protocol Number - *if Protocol = IP*
- Source Port - *if Protocol = TCP|UDP|SCTP*
- Destination Port - *if Protocol = TCP|UDP|SCTP*
- ICMP Type - *if Protocol = ICMPv6*
- ICMP Code - *if Protocol = ICMPv6*

Support NAT 64 CLAT (244986)

NAT64 CLAT traffic is now supported by the FortiGate. CLAT traffic comes from devices that use the SIIT translator that plays a part in affecting IPv6 - IPv4 NAT translation.

VIPs can contain FQDNs (268876)

Instead of mapping to an IP address VIP can use a Fully Qualified Domain Name. This has to be configured in the CLI and the FQDN must be an address object that is already configured in the address listing.

The syntax for using a FQDN is as follows:

```
config firewall vip
  edit <VIP id>
    set type fqdn
    set mapped-addr <FQDN address object>
  end
```

Access Control Lists (ACLs) (293399)

The access control list (ACL) feature allows you to deny IPv4 or IPv6 packets received at an NP6-accelerated interface based on source and destination address and service. If you add an access control policy to an interface, ACL checking is one of the first things that happens to the packet and checking is done by the NP6 processor. The result is very efficient protection that does not use CPU or memory resources.

In the GUI, the feature can be found at **Policy & Objects > IPv4 Access Control List Policy & Objects > IPv6 Access Control List**.

To add an IPv4 ACL through the CLI use the following syntax:

```
config firewall acl
  edit <acl Policy ID #>
    set status enable
    set interface <interface>
    set srcaddr <address object>
    set dstaddr <address object>
    set service <service object>
  end
end
```

To add an IPv6 ACL through the CLI use the following syntax:

```
config firewall acl6
  edit <acl Policy ID #>
    set status enable
    set interface <interface>
    set srcaddr <address object>
    set dstaddr <address object>
```

```
    set service <service object>
  end
end
```

GUI improvement for DoS Policy configuration (286905)

The user can now set the **Action**, whether **Pass** or **Block**, for all of the anomalies in a list at once when configuring a DoS policy. Just choose the desired option in the heading at the top of the column.

Expired Policy Object warnings (259338)

The Policy window indicates when a policy has become invalid due to its schedule parameters referring only to times in the past.

Firewall concepts

There are a number of foundational concepts that are necessary to have a grasp of before delving into the details of how the FortiGate firewall works. Some of these concepts are consistent throughout the firewall industry and some of them are specific to more advanced firewalls such as the FortiGate. Having a solid grasp of these ideas and terms can give you a better idea of what your FortiGate firewall is capable of and how it will be able to fit within your networks architecture.

This chapter describes the following firewall concepts:

- What is a Firewall?
- FortiGate Modes
- How Packets are handled by FortiOS
- Interfaces and Zones
- IPv6
- NAT
- Quality of Service

What is a Firewall?

The term firewall originally referred to a wall intended to confine a fire or potential fire within a building. Later uses refer to similar structures, such as the metal sheet separating the engine compartment of a vehicle or aircraft from the passenger compartment.

A firewall can either be software-based or hardware-based and is used to help keep a network secure. Its primary objective is to control the incoming and outgoing network traffic by analyzing the data packets and determining whether it should be allowed through or not, based on a predetermined rule set. A network's firewall builds a bridge between an internal network that is assumed to be secure and trusted, and another network, usually an external (inter)network, such as the Internet, that is not assumed to be secure and trusted.

Network Layer or Packet Filter Firewalls

Stateless Firewalls

Stateless firewalls are the oldest form of these firewalls. They are faster and simple in design requiring less memory because they process each packet individually and don't require the resources necessary to hold onto packets like stateful firewalls. Stateful firewalls inspect each packet individually and check to see if it matches a predetermined set of rules. According to the matching rule the packet is either be allowed, dropped or rejected. In the case of a rejection an error message is sent to the source of the traffic. Each packet is inspected in isolation and information is only gathered from the packet itself. Simply put, if the packets were not specifically allowed according to the list of rules held by the firewall they were not getting through.

Stateful Firewalls

Stateful firewalls retain packets in memory so that they can maintain context about active sessions and make judgments about the state of an incoming packet's connection. This enables Stateful firewalls to determine if a packet is the start of a new connection, a part of an existing connection, or not part of any connection. If a packet is part of an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing. If a packet does not match an existing connection, it will be evaluated according to the rules set for new connections. Predetermined rules are used in the same way as a stateless firewall but they can now work with the additional criteria of the state of the connection to the firewall.

Best Practices Tip for improving performance:



Blocking the packets in a denied session can take more cpu processing resources than passing the traffic through. By putting denied sessions in the session table, they can be kept track of in the same way that allowed session are so that the FortiGate unit does not have to redetermine whether or not to deny all of the packets of a session individually. If the session is denied all packets of that session are also denied.

In order to configure this you will need to use 2 CLI commands

```
config system setting
    set ses-denied-traffic enable
    set block-session-timer <integer 1 - 300> (this determines in
seconds how long, in seconds, the session is kept in the table)
end
```

Application Layer Firewalls

Application layer filtering is yet another approach and as the name implies it works primarily on the Application Layer of the OSI Model.

Application Layer Firewalls actually, for lack of a better term, understand certain applications and protocols. Examples would be FTP, DNS and HTTP. This form of filtration is able to check to see if the packets are actually behaving incorrectly or if the packets have been incorrectly formatted for the protocol that is indicated. This process also allows for the use of deep packet inspection and the sharing of functionality with Intrusion Prevention Systems (IPS).

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually dropping them without acknowledgment to the sender). Application firewalls work much like a packet filter but application filters apply filtering rules (allow/block) on a per process basis instead of filtering connections on a per port basis.

On inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. The additional inspection criteria can add extra latency to the forwarding of packets to their destination.

Proxy Servers

A proxy server is an appliance or application that acts as an intermediary for communicating between computers. A computer has a request for information. The packets are sent to the designated resource but before they can get there they are blocked by the proxy server saying that it will take the request and pass it on. The Proxy Server processes the request and if it is valid it passes onto the designated computer. The designated computer gets the packet and processes the request, sending the answer back to the proxy server. The proxy server sends the information back to the originating computer. It's all a little like a situation with two people who refuse to talk directly with each other using someone else to take messages back and forth.

From a security stand point a Proxy Server can serve a few purposes:

- Protects the anonymity of the originating computer
- The two computers never deal directly with each other
- Packets that are not configured to be forwarded are dropped before reaching the destination computer.
- If malicious code is sent it will affect the Proxy server with out affecting the originating or sending computer.

Proxies can perform a number of roles including:

- Content Filtering
- Caching
- DNS proxy
- Bypassing Filters and Censorship
- Logging and eavesdropping
- Gateways to private networks
- Accessing service anonymously

Security Profiles

Unified Threat Management and Next Generation Firewall are terms originally coined by market research firms and refer to the concept of a comprehensive security solution provided in a single package. It is basically combining of what used to be accomplished by a number of different security technologies all under a single umbrella or in this case, a single device. On the FortiGate firewall this is achieved by the use of Security Profiles and optimized hardware.

In effect it is going from a previous style of firewall that included among its features:

- Gateway Network Firewall
- Routing
- VPN

To a more complete system that includes:

- Gateway Network Firewall
- Routing
- VPN
- Traffic Optimization
- Proxy Services
- Content Filtering
- Application Control

- Intrusion Protection
- Denial of Service Attack Protection
- Anti-virus
- Anti-spam
- Data Leak Prevention
- Endpoint Control of Security Applications
- Load Balancing
- WiFi Access Management
- Authentication Integration into Gateway Security
- Logging
- Reporting

Advantages of using Security Profiles

- Avoidance of multiple installations.
- Hardware requirements are fewer.
- Fewer hardware maintenance requirements.
- Less space required.
- Compatibility - multiple installations of products increase the probability of incompatibility between systems.
- Easier support and management.
- There is only one product to learn therefore a reduced requirement of technical knowledge.
- Only a single vendor so there are fewer support contracts and Service Level Agreements.
- Easier to incorporated into existing security architecture.
- Plug and play architecture.
- Web based GUI for administration.

FortiGate Modes

The FortiGate unit has a choice of modes that it can be used in, either NAT/Route mode or Transparent mode. The FortiGate unit is able to operate as a firewall in both modes, but some of its features are limited in Transparent mode. It is always best to choose which mode you are going to be using at the beginning of the set up. Once you start configuring the device, if you want to change the mode you are going to lose all configuration settings in the change process.

NAT/Route Mode

NAT/Route mode is the most commonly used mode by a significant margin and is thus the default setting on the device. As the name implies the function of NAT is commonly used in this mode and is easily configured but there is no requirement to use NAT. The FortiGate unit performs network address translation before IP packets are sent to the destination network.

These are some of the characteristics of NAT/Route mode:

- Typically used when the FortiGate unit is a gateway between private and public networks.
- Can act as a router between multiple networks within a network infrastructure.

- When used, the FortiGate unit is visible to the networks that it is connected to.
- Each logical interface is on a distinct subnet.
- Each interface needs to be assigned a valid IP address for the subnet that it is connected to it.

Transparent Mode

Transparent mode is so named because the device is effectively transparent in that it does not appear on the network in the way that other network devices show as nodes in the path of network traffic. Transparent mode is typically used to apply the FortiOS features such as Security Profiles etc. on a private network where the FortiGate unit will be behind an existing firewall or router.

These are some of the characteristics of Transparent mode:

- The FortiGate unit is invisible to the network.
- All of its interfaces are on the same subnet and share the same IP address.
- The FortiGate unit uses a Management IP address for the purposes of Administration.
- Still able to use NAT to a degree, but the configuration is less straightforward

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools.

How Packets are handled by FortiOS

To give you an idea of what happens to a packet as it makes its way through the FortiGate unit, here is a brief overview. This particular trip of the packet is starting on the Internet side of the FortiGate firewall and ends with the packet exiting to the Internal network. An outbound trip would be similar. At any point in the path, if the packet is going through what would be considered a filtering process and if it fails the filter check, the packet is dropped and does not continue any further down the path.

This information is covered in more detail in other parts of the Troubleshooting chapter of the FortiOS Handbook in the Life of a Packet section.

The incoming packet arrives at the external interface. This process of entering the device is referred to as **ingress**.

Step #1 - Ingress

1. Denial of Service Sensor
2. IP integrity header checking
3. IPsec connection check
4. Destination NAT
5. Routing

Step #2 - Stateful Inspection Engine

1. Session Helpers
2. Management Traffic
3. SSL VPN
4. User Authentication

5. Traffic Shaping
6. Session Tracking
7. Policy lookup

Step #3 - Security Profiles scanning process

1. Flow-based Inspection Engine
2. IPS
3. Application Control
4. Data Leak Prevention
5. Email Filter
6. Web Filter
7. Anti-virus
8. Proxy-based Inspection Engine
9. VoIP Inspection
10. Data Leak Prevention
11. Email Filter
12. Web Filter
13. Anti-virus
14. ICAP

Step #4 - Egress

1. IPsec
2. Source NAT
3. Routing

Interfaces and Zones

A Firewall is a gateway device that may be the nexus point for more than 2 networks. The interface that the traffic is coming in on and should be going out on is a fundamental concern for the purposes of routing as well as security. Routing, policies and addresses are all associated with interfaces. The interface is essentially the connection point of a subnet to the FortiGate unit and once connected can be connected to other subnets.

Physical interfaces or not the only ones that need to be considered. There are also virtual interfaces that can be applied to security policies. VLANs are one such virtual interface. Interfaces if certain VPN tunnels are another.

Policies are the foundation of the traffic control in a firewall and the Interfaces and addressing is the foundation that policies are based upon. Using the identity of the interface that the traffic connects to the FortiGate unit tells the firewall the initial direction of the traffic. The direction of the traffic is one of the determining factors in deciding how the traffic should be dealt with. You can tell that interfaces are a fundamental part of the policies because, by default, this is the criteria that the policies are sorted by.

Zones are a mechanism that was created to help in the administration of the firewalls. If you have a FortiGate unit with a large number of ports and a large number of nodes in your network the chances are high that there is going to be some duplication of policies. Zones provide the option of logically grouping multiple virtual and physical FortiGate firewall interfaces. The zones can then be used to apply security policies to control the

incoming and outgoing traffic on those interfaces. This helps to keep the administration of the firewall simple and maintain consistency.

For example you may have several floors of people and each of the port interfaces could go to a separate floor where it connects to a switch controlling a different subnet. The people may be on different subnets but in terms of security they have the same requirements. If there were 4 floors and 4 interfaces a separate policy would have to be written for each floor to be allowed out on to the Internet off the WAN1 interface. This is not too bad if that is all that is being done, but now start adding the use of more complicated policy scenarios with Security Profiles, then throw in a number of Identity based issues and then add the complication that people in that organization tend to move around in that building between floors with their notebook computers.

Each time a policy is created for each of those floors there is a chance of an inconsistency cropping up. Rather than make up an additional duplicate set of policies for each floor, a zone can be created that combines multiple interfaces. And then a single policy can be created that uses that zone as one side of the traffic connection.

Access Control Lists

Access Control Lists (ACLs) in the FortiOS firmware could be considered a granular or more specifically targeted blacklist. These ACLs drop IPv4 or IPv6 packets at the physical network interface before the packets are analyzed by the CPU. On a busy appliance this can really help the performance.

The ACL feature is available on FortiGate with NP6-accelerated interfaces. ACL checking is one of the first things that happens to the packet and checking is done by the NP6 processor. The result is very efficient protection that does not use CPU or memory resources.

Incoming Interfaces

The configuration of the Access Control List allow you to specify which interface the ACL will be applied to. There is a hardware limitation that needs to be taken into account. The ACL is a Layer 2 function and is offloaded to the ISF hardware, therefore no CPU resources are used in the processing of the ACL. It is handled by the inside switch chip which can do hardware acceleration, increasing the performance of the FortiGate. The drawback is that the ACL function is only supported on switch fabric driven interfaces. It also cannot be applied to hardware switch interfaces or their members. Ports such as WAN1 or WAN2 that are found on some models that use network cards that connect to the CPU through a PCIe bus will not support ACL.

Addresses

Because the address portion of an entry is based on a FortiGate address object, it can be any of the address types used by the FortiGate, including address ranges. There is further granularity by specifying both the source and destination addresses. The traffic is blocked not on an either or basis of these addresses but the combination of the two, so that they both have to be correct for the traffic to be denied. Of course, if you want to block all of the traffic from a specific address all you have to do is make the destination address "all".

Because the blocking takes place at the interface based on the information in the packet header and before any processing such as NAT can take place, a slightly different approach may be required. For instance, if you are trying to protect a VIP which has an external address of x.x.x.x and is forwarded to an internal address of y.y.y.y, the destination address that should be used is x.x.x.x, because that is the address that will be in the packet's header when it hits the incoming interface.

Services

Further granulation of the filter by which the traffic will be denied is done by specifying which service the traffic will use.

IPv6

Internet Protocol version 6 (IPv6) will succeed IPv4 as the standard networking protocol of the Internet. IPv6 provides a number of advances over IPv4 but the primary reason for its replacing IPv4 is its limitation in addresses. IPv4 uses 32 bit addresses which means there is a theoretical limit of 2 to the power of 32. The IPv6 address scheme is based on a 128 bit address or a theoretical limit of 2 to the power of 128.

Possible Addresses:

- IPv4 = 4,294,967,296 (over 4 billion)
- IPv6 = 340,282,366,920,938,463,463,374,607,431,768,211,456 (over 340 undecillion - We had to look that term up. We didn't know what a number followed by 36 digits was either)

Assuming a world population of approximately 8 billion people, IPv6 would allow for each individual to have approximately 42,535,295,865,117,200,000,000,000,000 devices with an IP address. That's 42 quintillion devices.

There is little likelihood that you will ever need to worry about these numbers as any kind of serious limitation in addressing but they do give an idea of the scope of the difference in the available addressing.

Aside from the difference of possible addresses there is also the different formatting of the addresses that will need to be addressed.

A computer would view an IPv4 address as a 32 bit string of binary digits made up of 1s and 0s, broken up into 4 octets of 8 digits separated by a period "."

Example:

```
10101100.00010000.11111110.00000001
```

To make number more user friendly for humans we translate this into decimal, again 4 octets separated by a period "." which works out to:

```
172.16.254.1
```

A computer would view an IPv6 address as a 128 bit string of binary digits made up of 1s and 0s, broken up into 8 octets of 16 digits separated by a colon ":"

```
1000000000000001:0000110110111000:101011000001000:1111111000000001:0000000000000000
0:0000000000000000:0000000000000000:0000000000000000
```

To make number a little more user friendly for humans we translate this into hexadecimal, again 8 octets separated by a colon ":" which works out to:

```
8001:0DB8:AC10:FE01:0000:0000:0000:0000:
```

Because any four-digit group of zeros within an IPv6 address may be reduced to a single zero or altogether omitted, this address can be shortened further to:

```
8001:0DB8:AC10:FE01:0:0:0:0
```

or

8001:0DB8:AC10:FE01::

Some of the other benefits of IPv6 include:

- More efficient routing
- Reduced management requirement
- Stateless auto-reconfiguration of hosts
- Improved methods to change Internet Service Providers
- Better mobility support
- Multi-homing
- Security
- Scoped address: link-local, site-local and global address space

IPv6 in FortiOS

From an administrative point of view IPv6 works almost the same as IPv4 in FortiOS. The primary difference is the use IPv6 format for addresses. There is also no need for NAT if the FortiGate firewall is the interface between IPv6 networks. If the subnets attached to the FortiGate firewall are IPv6 and IPv4 NAT can be configured between the 2 different formats. This will involve either configuring a dual stack routing or IPv4 tunneling configuration. The reason for this is simple. NAT was developed primarily for the purpose of extending the number of usable IPv4 addresses. IPv6's addressing allows for enough available addresses so the NAT is no longer necessary.

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network. Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

By default the IPv6 settings are not displayed in the Web-based Manager. It is just a matter of enabling the display of these feature to use them through the web interface. To enable them just go to **System > Admin > Settings** and select **IPv6 Support on GUI**. Once enabled, you will be able to use IPv6 addresses as well as the IPv4 addressing for the following FortiGate firewall features:

- Static routing
- Policy Routing
- Packet and network sniffing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- IPsec VPN
- DNS
- DHCP
- SSL VPN
- Network interface addressing
- Security Profiles protection
- Routing access lists and prefix lists
- NAT/Route and Transparent mode

- NAT 64 and NAT 66
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Logging and reporting
- Security policies
- SNMP
- Authentication
- Virtual IPs and groups
- IPv6 over SCTP
- IPv6-specific troubleshooting, such as ping6

Dual Stack routing configuration

Dual stack routing implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. A dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary. Devices that are on this type of network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. In the FortiOS dual stack architecture it is not just the basic addressing functions that operate in both versions of IP. The other features of the appliance such as Security Profiles and routing can also use both IP stacks.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunneling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

IPv6 Tunneling

IPv6 Tunneling is the act of tunneling IPv6 packets from an IPv6 network through an IPv4 network to another IPv6 network. This is different than Network Address Translation (NAT) because once the packet reaches its final destination the true originating address of the sender will still be readable. The IPv6 packets are encapsulated within packets with IPv4 headers, which carry their IPv6 payload through the IPv4 network. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

The key to IPv6 tunneling is the ability of the 2 devices, whether they are a host or a network device, to be dual stack compatible. They have to be able to work with both IPv4 and IPv6 at the same time. In the process the entry node of the tunnel portion of the path will create an encapsulating IPv4 header and transmit the encapsulated packet. The exit node at the end of the tunnel receives the encapsulated packet. The IPv4 header is removed. The IPv6 header is updated and the IPv6 packet is processed.

There are two types of tunnels in IPv6:

Automatic tunnels	
	Automatic tunnels are configured by using IPv4 address information embedded in an IPv6 address – the IPv6 address of the destination host includes information about which IPv4 address the packet should be tunneled to.

Configured tunnels

Configured tunnels must be configured manually. These tunnels are used when using IPv6 addresses that do not have any embedded IPv4 information. The IPv6 and IPv4 addresses of the endpoints of the tunnel must be specified.

Tunnel Configurations

There are a few ways in which the tunneling can be performed depending on which segment of the path between the end points of the session the encapsulation takes place.

Network Device to Network Device	Dual stack capable devices connected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the path taken by the IPv6 packets.
Host to Network Device	Dual stack capable hosts can tunnel IPv6 packets to an intermediary IPv6 or IPv4 network device that is reachable through an IPv4 infrastructure. This type of tunnel spans the first segment of the path taken by the IPv6 packets.
Host to Host	Dual stack capable hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire path taken by the IPv6 packets.
Network Device to Host	Dual stack capable network devices can tunnel IPv6 packets to their final destination IPv6 or IPv4 host. This tunnel spans only the last segment of the path taken by the IPv6 packets.

Regardless of whether the tunnel starts at a host or a network device, the node that does the encapsulation needs to maintain soft state information, such as the maximum transmission unit (MTU), about each tunnel in order to process the IPv6 packets.

Tunneling IPv6 through IPsec VPN

A variation on the tunneling IPv6 through IPv4 is using an IPsec VPN tunnel between two FortiGate devices. FortiOS supports IPv6 over IPsec. In this sort of scenario, 2 networks using IPv6 behind FortiGate units are separated by the Internet, which uses IPv4. An IPsec VPN tunnel is created between the 2 FortiGate units and a tunnel is created over the IPv4 based Internet but the traffic in the tunnel is IPv6. This has the additional advantage of making the traffic secure as well.

NAT

NAT or Network Address Translation is the process that enables a single device such as a router or firewall to act as an agent between the Internet or Public Network and a local or private network. This “agent”, in real time, translates the source IP address of a device on one network interface, usually the Internal, to a different IP address as it leaves another interface, usually the interface connected to the ISP and the Internet. This enables a single public address to represent a significantly larger number of private addresses.

The Origins of NAT

In order to understand NAT it helps to know why it was created. At one time, every computer that was part of a network had to have its own addresses so that the other computers could talk to it. There were a few protocols in use at the time, some of which were only for use on a single network, but of those that were routable, the one that had become the standard for the Internet was IP (Internet Protocol) version 4.

When IP version 4 addressing was created nobody had any idea how many addresses would be needed. The total address range was based on the concept of 2 to the 32nd power, which works out to be 4 294 967 296 potential addresses. Once you eliminate some of those for reserved addresses, broadcast addresses, network addresses, multicasting, etc., you end up with a workable scope of about 3.2 million addressees. This was thought to be more than enough at the time. The designers were not expecting the explosion of personal computing, the World Wide Web or smart phones. As of the beginning of 2012, some estimate the number of computers in the world in the neighborhood of 1 billion, and most of those computer users are going to want to be on the Internet or Search the World Wide Web. In short, we ran out of addresses.

This problem of an address shortage was realized before we actually ran out, and in the mid 1990s 2 technical papers called RFCs numbered 1631 (<http://www.ietf.org/rfc/rfc1631.txt>) and 1918 (<http://tools.ietf.org/html/rfc1918>), proposed components of a method that would be used as a solution until a new addressing methodology could be implemented across the Internet infrastructure. For more information on this you can look up IP version 6.

RFC 1631 described a process that would allow networking devices to translate a single public address to multiple private IP addresses and RFC 1918 laid out the use of the private addresses. The addresses that were on the Internet (Public IP addresses) could not be duplicated for them to work as unique addresses, but behind a firewall, which most large institutions had, they could use their own Private IP addresses for internal use and the internal computers could share the external or Public IP address.

To give an idea on a small scale how this works, image that a company has a need for 200 computer addresses. Before Private IP addresses and NAT the company would have purchased a full Class C address range which would have been 254 usable IP addresses; wasting about 50 addresses. Now with NAT, that company only needs 1 IP address for its 200 computers and this leaves the rest of the IP addresses in that range available for other companies to do the same thing.

NAT gives better value than it would first appear because it is not 253 companies that can use 254 addresses but each of those 254 companies could set up their networking infrastructures to use up to thousands of Private IP addresses, more if they don't all have to talk to the Internet at the same time. This process enabled the Internet to keep growing even though we technically have many more computers networked than we have addresses.

Dynamic NAT

Dynamic NAT maps the private IP addresses to the first available Public Address from a pool of possible Addresses. In the FortiGate firewall this can be done by using IP Pools.

Overloading

This is a form of Dynamic NAT that maps multiple private IP address to a single Public IP address but differentiates them by using a different port assignment. This is probably the most widely used version of NAT. This is also referred to as PAT (Port Address Translation) or Masquerading.

An example would be if you had a single IP address assigned to you by your ISP but had 50 or 60 computers on your local network.

Say the internal address of the interface connected to the ISP was 256.16.32.65 (again an impossible address) with 256.16.32.64 being the remote gateway. If you are using this form of NAT any time one of your computers accesses the Internet it will be seen from the Internet as 256.16.32.65. If you wish to test this go to 2 different computers and verify that they each have a different private IP address then go to a site that tells you your IP address such as www.ipchicken.com. You will see that the site gives the same result of 256.16.32.65, if it existed, as the public address for both computers.

As mentioned before this is sometimes called Port Address Translation because network device uses TCP ports to determine which internal IP address is associated with each session through the network device. For example, if you have a network with internal addresses ranging from 192.168.1.1 to 192.168.1.255 and you have 5 computers all trying to connect to a web site which is normally listening on port 80 all of them will appear to the remote web site to have the IP address of 256.16.32.65 but they will each have a different sending TCP port, with the port numbers being somewhere between 1 and 65 535, although the port numbers between 1 to 1024 are usually reserved or already in use. So it could be something like the following:

192.168.1.10	256.16.32.65:	port 486
192.168.1.23	256.16.32.65:	port 2409
192.168.1.56	256.16.32.65:	port 53763
192.168.1.109	256.16.32.65:	port 5548
192.168.1.201	256.16.32.65:	port 4396

And the remote web server would send the responding traffic back based on those port numbers so the network device would be able to sort through the incoming traffic and pass it on to the correct computer.

Overlapping

Because everybody is using the relative same small selection of Private IP addresses it is inevitable that there will be two networks that share the same network range that will need to talk with each other. This happens most often over Virtual Private Networks or when one organization ends up merging with another. This is a case where a private IP address may be translated into a different private IP address so there are no issues with conflict of addresses or confusion in terms of routing.

An example of this would be when you have a Main office that is using an IP range of 172.16.0.1 to 172.20.255.255 connecting through a VPN to a recently acquired branch office that is already running with an IP range of 172.17.1.1 to 172.17.255.255. Both of these ranges are perfectly valid but because the Branch office range is included in the Main Office range any time the system from the Main office try to connect to an address in the Branch Office the routing the system will not send the packet to the default gateway because according to the routing table the address is in its own subnet.

The plan here would be to NAT in both directions so that traffic from neither side of the firewall would be in conflict and they would be able to route the traffic. Everything coming from the Branch Office could be assigned an address in the 192.168.1.1 to 192.168.1.255 range and everything from the Main office going to the Branch Office could be assigned to an address in the 192.168.10.1 to 192.168.10.255 range.

Static NAT

In Static NAT one internal IP address is always mapped to the same public IP address.

In FortiGate firewall configurations this is most commonly done with the use of Virtual IP addressing.

An example would be if you had a small range of IP addresses assigned to you by your ISP and you wished to use one of those IP address exclusively for a particular server such as an email server.

Say the internal address of the Email server was 192.168.12.25 and the Public IP address from your assigned addresses range from 256.16.32.65 to 256.16.32.127. Many readers will notice that because one of the numbers

is above 255 that this is not a real Public IP address. The Address that you have assigned to the interface connected to your ISP is 256.16.32.66, with 256.16.32.65 being the remote gateway. You wish to use the address of 256.16.32.70 exclusively for your email server.

When using a Virtual IP address you set the external IP address of 256.16.32.70 to map to 192.168.12.25. This means that any traffic being sent to the public address of 256.16.32.70 will be directed to the internal computer at the address of 192.168.12.25

When using a Virtual IP address, this will have the added function that when ever traffic goes from 192.168.12.25 to the Internet it will appear to the recipient of that traffic at the other end as coming from 256.16.32.70.

You should note that if you use Virtual IP addressing with the Port Forwarding enabled you do not get this reciprocal effect and must use IP pools to make sure that the outbound traffic uses the specified IP address.

Benefits of NAT

More IP addresses Available while Conserving Public IP Addresses

As explained earlier, this was the original intent of the technology and does not need to be gone into further.

Financial Savings

Because an organization does not have to purchase IP addresses for every computer in use there is a significant cost savings due to using the process of Network Address Translation.

Security Enhancements

One of the side benefits of the process of NAT is an improvement in security. Individual computers are harder to target from the outside and if port forwarding is being used computers on the inside of a firewall are less likely to have unmonitored open ports accessible from the Internet.

Ease of Compartmentalization of Your Network

With a large available pool of IP addresses to use internally a network administrator can arrange things to be compartmentalized in a rational and easily remembered fashion and networks can be broken apart easily to isolate for reasons of network performance and security.

Example

You have a large organization that for security reasons has certain departments that do not share network resources.

You can have the main section of the organization set up as follows;

Network Devices	192.168.1.1 to 192.168.1.25
Internal Servers	192.168.1.26 to 192.168.1.50
Printers	192.168.1.51 to 192.168.1.75

Administration Personnel	192.168.1.76 to 192.168.1.100
Sales People	192.168.1.101 to 192.168.1.200
Marketing	192.168.1.201 to 192.168.1.250

You could then have the following groups broken off into separate subnets:

Accounting	192.168.100.1 to 192.168.100.255
Research and Development	172.16.1.1 to 172.16.255.255
Executive Management	192.168.50.1 to 192.168.50.255
Web sites and Email Servers	10.0.50.1 to 10.0.50.255

These addresses do not have to be assigned right away but can be used as planned ranges.

NAT in Transparent Mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.
- Add IP pools as required for source address translation

A FortiGate unit operating in Transparent mode normally has only one IP address - the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

Use the following steps to configure NAT in Transparent mode:

1. Add two management IPs
2. Add an IP pool to the WAN1 interface
3. Add an Internal to WAN1 security policy

You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

The usual practice of NATing in transparent mode makes use of two management IP addresses that are on different subnets, but this is not an essential requirement in every case.

If there is a router between the client systems and the FortiGate unit you can use the router's capabilities of tracking sessions to assign NATed addresses from an IP pool to the clients even if the assigned address don't belong to a subnet on your network.

Example

Client computer has an IP address of 1.1.1.33 on the subnet 1.1.1.0/24.

Router “A” sits between the client computer and the FortiGate (in Transparent mode) with the IP address of 1.1.1.1 on the client’s side of the router and the IP address of 192.168.1.211 on the FortiGate’s side of the router.

Use NAT to assign addresses from an address pool of 9.9.9.1 to 9.9.9.99 to traffic coming from gateway of 192.168.1.211.

To enable the return traffic to get to the original computer, set up a static route that assigns any traffic with a destination of 9.9.9.0/24 to go through the 192.168.1.211 gateway. As long as the session for the outgoing traffic has been maintained, communication between the client computer and the external system on the other side of the FortiGate will work.

Central NAT Table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well. NAT policies are applied to network traffic after a security policy.

NAT 64 and NAT46

NAT64 and NAT46 are the terms used to refer to the mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and vice-versa. Without such a mechanism an IPv6 node on a network such as a corporate LAN would not be able to communicate with a web site that was still in a IPv4 only environment and IPv4 environments would not be able to connect to IPv6 networks.

One of these setups involves having at least 2 interfaces, 1 on an IPv4 network and 1 on an IPv6 network. The NAT64 server synthesizes AAAA records, used by IPv6 from A records used by IPv4. This way client-server and peer to peer communications will be able to work between an IPv6 only client and an IPv4 server without making changes to either of the end nodes in the communication transaction. The IPv6 network attached to the FortiGate unit should be a 32 bit segment, (for instance 64:ff9b::/96, see RFC 6052 and RFC 6146). IPv4 address will be embedded into the communications from the IPv6 client.

Because the IPv6 range of addresses is so much larger than the IPv4 range, a one to one mapping is not feasible. Therefore the NAT64 function is required to maintain any IPv6 to IPv4 mappings that it synthesizes. This can be done either statically by the administrator or automatically by the service as the packets from the IPv6 network go through the device. The first method would be a stateless translation and the second would be a stateful translation. NAT64 is designed for communication initiated from IPv6 hosts to IPv4 addresses. It is address mapping like this that allows the reverse to occur between established connections. The stateless or manual

method is an appropriate solution when the NAT64 translation is taking place in front of legacy IPv4 servers to allow those specific servers to be accessed by remote IPv6-only clients. The stateful or automatic solution is best used closer to the client side when you have to allow some specific IPv6 clients to talk to any of the IPv4-only servers on the Internet.

There are currently issues with NAT64 not being able to make everything accessible. Examples would be SIP, Skype, MSN, Goggle talk, and sites with IPv4 literals. IPv4 literals being IPv4 addresses that are imbedded into content rather than a FQDN.

Policies that employ NAT64 or NAT46 can be configured from the web-based manager as long as the feature is enabled using the Features setting found at **System > Config > Features**.

- To create a NAT64 policy go to **Policy > Policy > NAT64 Policy** and select **Create New**.
- To create a NAT46 policy go to **Policy > Policy > NAT46 Policy** and select **Create New**.

The difference between these NAT policies and regular policies is that there is no option to use the security profiles and sensors.

NAT 66

NAT 66 is Network Address Translation between 2 IPv6 network. The basic idea behind NAT 66 is no different than the regular NAT between IPv4 networks that we are all used to. The difference are in the mechanics of how it is performed, mainly because of the complexity and size of the addresses that are being dealt with.

In an IPv4 world, the reason for the use of NAT was usually one or a combination of the following 3 reasons:

- Improved security - actual addresses behind NAT are virtually hidden
- Amplification of addresses - hundreds of computers can use as little as a single public IP address
- Internal address stability - there is control of internal addressing. The addresses can stay the same even if Internet Service Providers change.

In these days of security awareness the protective properties of NAT are not something that are not normally depended on by themselves to defend a network and with the vastly enlarged IPv6 address scope there is no longer a need to amplify the available addresses. However, the desire to have internal address control still exists. The most common reason for using NAT66 is likely to be the maintaining of the existing address scheme of the internal network despite changes outside of it. Imagine that you have an internal network of 2000 IP addresses and one day the company changes its ISP and thus the addresses assigned to it. Even if most of the addressing is handled by DHCP, changing the address scheme is going to have an impact on operations.

Addressing stability can be achieved by:

- Keeping the same provider - this would depend on the reason for the change. If the cost of this provider has become too expensive this is unlikely. If the ISP is out of business it becomes impossible.
- Transfer the addresses from the old provider to the new one - There is little motivation for an ISP to do you a favor for not doing business with them.
- Get your own autonomous system number - this can be too expensive for smaller organizations.
- NAT - this is the only one on the list that is in the control of IT.

There are differences between NAT66 and IPv4 NAT. Because there is no shortage of addresses most organizations will be given a /48 network that can be translated into another /48 network. This allows for a one to one translation, no need for port forwarding. This is a good thing because port forwarding is more complicated in IPv6. In fact, NAT66 will actually just be the rewriting of the prefix on the address.

Example

If your current IPv6 address is

```
2001:db8:cafe::/48
```

you could change it to

```
2001:db8:fea7::/48
```

There is an exception to the one to one translation. NAT66 cannot translate internal networks that contain 0xffff in bits 49 through 63 - this is due to the way checksums are calculated in TCP/IP: they use the one's-complement representation of numbers which assigns the value zero to both 0x0000 and 0xffff.

How FortiOS differentiates sessions when NATing

The basics of NAT are fairly simple. Many private addresses get translated into a smaller number of public addresses, often just one. The trick is how the FortiGate keeps track of the return traffic because the web server, or what ever device that was out on the Internet is going to be sending traffic back not to the private address behind the FortiGate but to the IP address of the interface on the public side of the FortiGate.

The way this is done is by making each session unique. Most of the attributes that are available in the network packets cannot be changed without changing where the packet will go but because the source port has to be changed anyway in case two computer on the network used the same source port this is a useful way of making each listing of network attributes a unique combination. As a packet goes through the NAT process FortiOS assigns different source ports for each of the internally initiated sessions and keeping track of which port was used for each device in a database until the session has ended. It then becomes a matter of how the port number is selected.

In a very simple example of an environment using NAT, we will use a fictitious university with a rather large student population. So large in fact that they use a subnet of 10.0.0.0/8 as their subnet for workstation IP addresses. All of these private IP addresses are NATed out a single IP address. To keep the number of numeric values in this example from getting to a confusing level, we'll just use "u.u.u.1" to refer to the public IP address of the University and the IP address of the web server on the Internet will be "w.w.w.1".

Student A (IP address 10.1.1.56) sends an HTML request to a web server on the Internet with the IP address w.w.w.1. The applicable networking information in the packet breaks down as follows:

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

The source IP address is now that of the public facing interface of the FortiGate and source port number is an unused TCP port number on the FortiGate chosen by the FortiGate. Of these variable the only one the that FortiGate can really change and still have the packet reach the correct destination, in both directions, is the source port number.

There are a few methods of assigning the port number. First we'll look at the methods that are or have been used in the industry but aren't used by Fortinet.

Global pool

This method of differentiation focuses on the attribute of the source port number. In this approach a single pool of potential port numbers is set aside for the purposes of NAT. As a pool number is assigned, it is removed from the pool so that two sessions from different computers can not using the same port number. Once the session is over and no longer in use by the computer, the port number is put back into the pool where it can be assigned again.

Example global pool:

	Hexidecimal	Decimal
Start or range	0x7000	28672
End end of range	0xF000	61440
Possible ports in range	215	32768

This is a simple approach to implement and is good if the number of connections in unlike to reach the pool size. It would be okay for home use, but our example is for a university using $10.1.1.0/8$ as a subnet. That means 16,777,214 possible IP addresses; more than this method can handle.

Fortinet does not use this method.

Global per protocol

This method uses the attributes source port number and type of protocol to differentiate between sessions. This approach is a variation of the first one. An additional piece of information is referred to in the packet that describes the protocol. For instance UDP or TCP. This could effectively double the number of potential addresses to NAT.

Example:

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	udp	udp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

Even though the source port is the same, because the protocol is different they are considered to be from different sessions and different computers.

The drawback is that it would depend on the protocols being used be evenly distributed between TCP and UDP. Even if this was the case the number would only double; reaching an upper limit of 65,536 possible connections. That number is still far short of the possible more than 16 million for an IP subnet with an eight bit subnet mask like the one in our example.

Fortinet does not use this method.

Per NAT IP Pool

This approach adds on to the previous one by adding another variable. In this case that variable is the IP addresses on the public side of the FortiGate. By having a pool of IP addresses to assign as the source IP address when NATing, the same number that was potentially available for the Global per protocol method can be multiplied by the number of external IP addresses in the pool. If you can assign a second IP address to the pool, you can double the potential number of sessions.

Example:

In this example it will be assumed that the FortiGate has 2 IP addresses that it can use. This could happen either by using two ISPs, or by having a pool of IP addresses assigned to a single interface. For simplicity will refer to these IP public IP addresses as u.u.u.1 and u.u.u.2.

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp

Attribute	Original Packet	Packet after NATing
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.2
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

In this example we even made the protocol the same. After the NATing process all of the variables are the same except the source address. This is still going to make it back to the original sender.

The drawback is that if you have only one IP address for the purposes of NATing this method does not gain you anything over the last method. Or if you do have multiple IP addresses to use it will still take quite a few to reach the 16 million possible that the subnet is capable of handling.

Fortinet does not use this method.

Per NAT IP, destination IP, port, and protocol

This is the approach that FortiOS uses.

It uses all of the differentiation point of the previous methods, NAT IP, port number and protocol, but the additional information point of the destination IP is also used. So now the network information points in the packet that the FortiGate keeps in its database to differentiate between sessions is:

- Public IP address of the FortiGate assigned by NATing
- Protocol of the traffic
- Source port assigned by the FortiGate
- Destination IP address of the packet

The last one is an especially good way to differentiate because as a theoretical number, the upper limit on that is the numbers of Public IP addresses on the whole of the Internet. Chances are that while a large number of session from inside the University will be going to a small group of sites such as Google, Youtube, Facebook and some others it is unlikely that they will all be going to them at the same time.

Example:

In this example it will be assumed that the FortiGate has only one IP address. Two possible packets will be described. The only difference in the attributes recorded will be the destination of the HTML request. These packets are still considered to be from different sessions and any responses will make it back to the correct computer.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.2	w.w.w.2
Protocol	tcp	tcp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

The reason that these attributes are used to determine differentiation between traffic is based on how the indexes for the sessions are recorded in the database. When a TCP connection is made through a FortiGate unit, a session is created and two indexes are created for the session. The FortiGate unit uses these indexes to guide matching traffic to the session.

This following could be the session record for the TCP connection in the first example.

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1

Attribute	Outgoing Traffic	Returning Traffic
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	80
Destination port	80	46372 (FortiGate assigned port)

Using the FortiGate's approach for session differentiation, FortiOS only has to ensure that the assigned port, along with the other four attributes is a unique combination to identify the session. So for example, if Student A simultaneously makes a HTTP(port 80) connection and a HTTPS(port 443) connection the same web server this would create another session and the index in the reply direction would be:

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	443
Destination port	443	46372 (FortiGate assigned port)

These two sessions are different and acceptable because of the different source port numbers on the returning traffic or the destination port depending on the direction of the traffic.

Calculations for possible session numbers

The result of using these four attributes instead of just the one that was originally used is a large increase in the number of possible unique combinations. For those who love math, the maximum number of simultaneous connections that can be supported is:

$$N \times R \times P \times D \times Dp$$

where:

- **N** is the number of NAT IP addresses
- **R** is the port range,
- **P** is the number of protocols,
- **D** is the number of unique destination IP addresses
- **Dp** the number of unique destination ports.

As a rough example let's do some basic calculations

- N - In our existing example we have already stated that there is only one public IP address that is being used by NAT. Realistically, for a university this number would likely be larger, but we're keeping it simple.

$$N = 1$$

R - The port range for our example has already been describe and we will keep it the same.

$$R = 32768$$

P - While there are a few protocols that are involved in Internet traffic we will limit this calculation just to TCP traffic.

$$P = 1$$

D - As mentioned before the number of unique destination addresses is growing larger every day, so figuring out the upper limit of that number would be difficult to say the least. Instead we will make the assumption that most of the university students, do to their shared interest and similar demographic will concentrate most of their web browsing to the same sites; sites such as YouTube, Facebook, Google, Twitter, Instagram, Wikipedia etc. This is not even taking into account the fact that many of these popular sites use load balancing and multiple IP addresses. As an arbitrary number let's use the number 25.

$$D = 25$$

Dp - To keep things simple it is tempting to limit the destination port to port 80, the one that many associate with web browsing, but this would not be realistic. the use of HTTPS, port 443 is on the rise. There is also email, DNS, FTP, NTP and a number of other background services that we use without thinking too closely about. Let's keep it small and say ten of them.

$$Dp = 10$$

The math on this very conservative calculation is:

$$1 \times 32768 \times 1 \times 25 \times 10 = 8,192,000 \text{ possible NAT sessions}$$

When you take into account that the chances of everybody being online at the same time, going only to one of those 25 sites and not millions of others, and using only TCP not UDP or any of the other protocols, it starts to look like this method may provide enough potential unique sessions even for a subnet as large as the one described.

IP Pools

IP Pools are a mechanism that allow sessions leaving the FortiGate Firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses will be used instead of the IP address assigned to that FortiGate interface.



When using IP pools for NATing, there is a limitation that must be taken into account. In order for communication to be successful in both directions, it is normal for the source address in the packet header assigned by the NAT process to be an address that is associated with the interface that the traffic is going through. For example, if traffic is going out an interface with the IP address 172.16.100.1, packets would be NATed so that the source IP address would be 172.16.100.1. This way the returning traffic will be directed to the same interface on the same FortiGate that the traffic left from. Even if the packets are assigned a source address that is associated with another interface on the same FortiGate this can cause issues with asymmetrical routing. It is possible to configure the NATed source IP address to be different than the IP address of the interface but you have to make sure that the routing rules of the surrounding network devices take this unorthodox approach into consideration.

There are 4 types of IP Pools that can be configured on the FortiGate firewall:

- One-to-One - in this case the only internal address used by the external address is the internal address that it is mapped to.
- Overload - this is the default setting. Internal addresses other than the one designated in the policy can use this address for the purposes of NAT.
- Fixed Port Range - rather than a single address to be used, there is a range of addresses that can be used as the NAT address. These addresses are randomly assigned as the connections are made.
- Port Block Allocation - this setting is used to allocate a block of port numbers for IP pool users. Two variables will also have to be set. The block size can be set from 64 to 4096 and as the name implies describes the number of ports in one block of port numbers. The number of blocks per user determines how many of these blocks will be assigned. This number can range from 1 to 128.



Be careful when calculating the values of the variables. The maximum number of ports that are available on an address is 65,536. If you chose the maximum value for both variables you will get a number far in excess of the available port numbers.

$$4096 \times 128 = 524,288$$

One of the more common examples is when you have an email server behind your FortiGate firewall and the range of IP addresses assigned to you by your ISP is more than one. If an organization is assigned multiple IP addresses it is normally considered a best practice to assign a specific address other than the one used for the Firewall to the mail server. However, when normal NAT is used the address assigned to the firewall is also assigned to any outbound sessions. Anti-spam services match the source IP address of mail traffic that they receive to the MX record on DNS servers as an indicator for spam. If there is a mismatch the mail may not get through so there is a need to make sure that the NATed address assigned matches the MX record.

You can also use the Central NAT table as a way to configure IP pools.

Source IP address and IP pool address matching when using a range

When the source addresses are translated to an IP pool that is a range of addresses, one of the following three cases may occur:

Scenario 1:

The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable fixed port in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Scenario 2:

The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable fixed port in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Scenario 3:

The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

ARP Replies

If a FortiGate firewall interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

$(1.1.1.0-1.1.1.255) \text{ and } (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20$

The port2 interface overlap IP range with IP_pool_2 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20$

The port2 interface overlap IP range with IP_pool_3 is:

$$(2.2.2.0-2.2.2.255) \& (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40$$

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select Enable NAT in a security policy and then select Dynamic IP Pool. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool. Whether or not the external address of an IP Pool will respond to an ARP request can be disabled. You might want to disable the ability to respond to ARP requests so that these address cannot be used as a way into your network or show up on a port scan.

IP pools and zones

Because IP pools are associated with individual interfaces IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

Fixed Port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

However, enabling the use of a fixed port means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Match-VIP

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. By default, the feature is disabled.

Services and TCP ports

There are a number of different services and protocols in use on the Internet. The most commonly known is HTTP which is used by web servers to transmit requests and responses for unencrypted web pages. These

services are set up to listen for requests on a numbered port. These services and protocols can use any port from 1 to 65,535. To keep things simple for everyone a large number of the more commonly used services started using a standardized list of ports. For instance, though it is not required, by default, most web servers listen for HTTP requests on port 80 and by default, web browsers will send HTTP traffic to port 80. If you wish to use another port such as 8080 you would put “:8080” at the end of the URL to indicate that you want the browser to use 8080 instead of the default port.

Example

Default URL for HTTP traffic when the web server is listening on the standard HTTP port:

`http://fortinet.com`

URL to the same address when the web server is listening for HTTP traffic on port 8080

`http://fortinet.com:8080`

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined on the FortiGate unit. If there is a service that does not appear on the list you can create a service or edit an existing one. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create a service.

Best Practices



While you can edit a predefined service it is best to leave those ones alone and create a new service and name it something similar such as the same service name with a descriptive identifier appended.

Based on the previous example, instead of the name “HTTP” you could name the service “HTTP8080” or use the application that is using that port, “HTTP-Application”.

Protocol Types

One of the fundamental aspects of a service is the type of protocol that use used to define it. When a service is defined one of the following categories of protocol needs to be determined:

- TCP/UDP/SCTP
- ICMP
- ICMP6
- IP

Depending on which of these protocol categories is choose another set of specifications will can also be defined.

TCP/UDP/SCTP

This is the most commonly used service protocol category. Once this category has been selected the other available options to choose are an address, either IP or FQDN, and the protocol and port number.

The protocol will be TCP, UDP or SCTP.

ICMP or ICMP6

When ICMP or ICMP6 is chosen the available options are the ICMP Type and its code.

IP

When IP is the chosen protocol type the addition option is the Protocol Number.

TCP

Transmission Control Protocol (TCP) is one of the core or fundamental protocols of the Internet. It is part of the Transport Layer of the OSI Model. It is designed to provide reliable delivery of data from a program on one device on the network or Internet to another program on another device on the network or Internet. TCP achieves its reliability because it is a connection based protocol. TCP is stream-oriented. It transports streams of data reliably and in order.

TCP establishes a prior connection link between the hosts before sending data. This is often referred to as the handshake. Once the link is established the protocol uses checks to verify that the data transmitted. If an error check fails the data is retransmitted. This makes sure that the data is getting to the destination error free and in the correct order so that it can be put back together into a form that is identical to the way they were sent.

TCP is configured more for reliability than for speed and because of this TCP will likely be slower than a connectionless protocol such as UDP. This is why TCP is generally not used for real time applications such as voice communication or online gaming.

Some of the applications that use TCP are:

- World Wide Web (HTTP and HTTPS)
- Email (SMTP, POP3, IMAP4)
- Remote administration (RDP)
- File transfer (FTP)

UDP

User Datagram Protocol (UDP) like TCP is one of the core protocols of the Internet and part of the Transport Layer of the OSI Model. UDP is designed more for speed than reliability and is generally used for different applications than TCP. UDP sends messages, referred to as datagrams across the network or Internet to other hosts without establishing a prior communication link. In other words, there is no handshake.

UDP is an unreliable service as the datagrams can arrive out of order, duplicated or go missing without any mechanism to verify them. UDP works on the assumption that any error checking is done by the application or is not necessary for the function of the application. This way it avoids the overhead that is required to verify the integrity of the data.

This lack of overhead improves the speed of the data transfer and is why UDP is often used by applications that are time sensitive in nature. UDP's stateless nature is also great for applications that answer a large number of small queries from a large number of clients.

Common uses for UDP are:

- Domain Name Resolution (DNS)
- Time (NTP)

- Streaming media (RTSP, RTP and RTCP)
- Telephone of the Internet (VoIP)
- File Transfer (TFTP)
- Logging (SNMP)
- Online games (GTP and OGP)

SCTP

Stream Control Transmission Protocol (SCTP) is part of the Transport Layer of the OSI Model just like TCP and UDP and provides some of the features of both of those protocols. It is message or datagram orientated like UDP but it also ensures reliable sequential transport of data with congestion control like TCP.

SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- Network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP uses multi-streaming to transport its messages which means that there can be several independent streams of messages traveling in parallel between the points of the transmission. The data is sent out in larger chunks of data than is used by TCP just like UDP but the messages include a sequence number within each message in the same way that TCP does so that the data can be reassembled at the other end of the transmission in the correct sequence without the data having to arrive in the correct sequence.

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a much newer protocol. It was defined by the IETF Signaling Transport (SIGTRAN) working group in 2000. It was introduced by RFC 3286 and more fully define by RFC 4960.

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to "ALL". FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism
- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Protocol Port Values

The source and destination ports for TCP/UDP/SCTP services are important to get correct. If they are reversed the service will not work. The destination port(s) are the ones that refer to the ports that the computer will be listening on. These are the port numbers that most people are familiar with when they associate a port number to a protocol. In most cases the source port will be one that is randomly assigned by the computer that is not being already used by another service.

Most people associate HTTP with port 80. This means that a web-server will be listening on port 80 for any http requests being sent to the computer. The computer that is sending the request can use any port that is not already assigned to another service or communication session. There are 65,535 ports that it can randomly assign, but because the ports from 1 to 1024 are normally used for listening for incoming communications it is usually not in that range. It is unless there is a specific instance when you know that a communication will be coming from a predefined source port it is best practice to set the source port range from 1 to 65,535.

ICMP

The Internet Control Message Protocol (ICMP) is a protocol layered onto the Internet Protocol Suite to provide error reporting flow control and first-hop gateway redirection. It is normally used by the operating systems of networked computers to send connectivity status query, response and error messages. It is assigned protocol number 1. There is a version of the protocol for both IPv4 and for IPv6. It is not designed to be absolutely reliable like TCP.

ICMP is not typically used for transporting data or for end-user network applications with the exception of some diagnostic utilities such as ping and traceroute.

ICMP messages are sent in several situations, for example:

- when a datagram cannot reach its destination,
- time exceeded messages
- redirect messages
- when the gateway does not have the buffering capacity to forward a datagram
- when the gateway can direct the host to send traffic on a shorter route.

Some of the specific ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

ICMP Types and Codes

ICMP has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

ICMP Types and Codes

Type Number	Type Name	Optional Code(s)
0	Echo Reply	
1	Unassigned	
2	Unassigned	

Type Number	Type Name	Optional Code(s)
3	Destination Unreachable	0 Net Unreachable
		1 Host Unreachable
		2 Protocol Unreachable
		3 Port Unreachable
		4 Fragmentation Needed and Don't Fragment was Set
		5 Source Route Failed
		6 Destination Network Unknown
		7 Destination Host Unknown
		8 Source Host Isolated
		9 Communication with Destination Network is Administratively Prohibited
		10 Communication with Destination Host is Administratively Prohibited
		11 Destination Network Unreachable for Type of Service
		12 Destination Host Unreachable for Type of Service
		13 Communication Administratively Prohibited
		14 Host Precedence Violation
		15 Precedence cutoff in effect
4	Source Quench	
5	Redirect	0 Redirect Datagram for the Network (or subnet)
		1 Redirect Datagram for the Host
		2 Redirect Datagram for the Type of Service and Network
		3 Redirect Datagram for the Type of Service and Host
6	Alternate Host Address	
7	Unassigned	

Type Number	Type Name	Optional Code(s)
8	Echo	
9	Router Advertisement	
10	Router Selection	
11	Time Exceeded	0 Time to Live exceeded in Transit 1 Fragment Reassembly Time Exceeded
12	Parameter Problem	0 Pointer indicates the error 1 Missing a Required Option 2 Bad Length
13	Timestamp	
14	Timestamp Reply	
15	Information Request	
16	Information Reply	
17	Address Mask Request	
18	Address Mask Reply	
19	Reserved (for Security)	
20 - 29	Reserved (for Robustness Experiment)	
30	Traceroute	
31	Datagram Conversion Error	
32	Mobile Host Redirect	

Type Number	Type Name	Optional Code(s)
33	IPv6 Where-Are-You	
34	IPv6 I-Am-Here	
35	Mobile Registration	
36	Mobile Registration Reply	
37	Domain Name Request	
38	Domain Name Reply	
39	SKIP	
40	Photuris	
41 - 255	Reserved	

log-invalid-packet

The `log-invalid-packet` CLI setting is one that is intended to log invalid ICMP packets. The exact definition being:

If the FortiGate unit receives an ICMP error packet that contains an embedded IP (A,B) | TCP (C,D) header, then if FortiOS can locate the A:C -> B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped.

When this field is enabled, the FortiGate also log messages that are not ICMP error packets.

Types of logs covered by log-invalid-packet

- Invalid ICMP
 - If ICMP error message verification (see "check-reset-range") is enabled
- Invalid DNS packets
 - DNS packets that contain requests for non-existing domains
- iprope check failed
- reverse path check fail
- denied and broadcast traffic
- no session matched

Some other examples of messages that are not errors that will be logged, based on [RFC792](#):

Type 3 messages correspond to "Destination Unreachable Message"

- Type 3, Code 1 = host unreachable
- Type 3, Code 3 = port unreachable

Type 11 messages correspond to "Time Exceeded Message"

- Type 11, Code 0 = time to live exceeded in transit

ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) is the new implementation of the Internet Control Message Protocol (ICMP) that is part of Internet Protocol version 6 (IPv6). The ICMPv6 protocol is defined in [RFC 4443](#).

ICMPv6 is a multipurpose protocol. It performs such things as:

- error reporting in packet processing
- diagnostic functions
- Neighbor Discovery process
- IPv6 multicast membership reporting

It also designed as a framework to use extensions for use with future implementations and changes.

Examples of extensions that have already been written for ICMPv6:

- Neighbor Discovery Protocol (NDP) - a node discovery protocol in IPv6 which replaces and enhances functions of ARP.
- Secure Neighbor Discovery Protocol (SEND) - an extension of NDP with extra security.
- Multicast Router Discovery (MRD) - allows discovery of multicast routers.

ICMPv6 messages use IPv6 packets for transportation and can include IPv6 extension headers. ICMPv6 includes some of the functionality that in IPv4 was distributed among protocols such as ICMPv4, ARP (Address Resolution Protocol), and IGMP (Internet Group Membership Protocol version 3).

ICMPv6 has simplified the communication process by eliminating obsolete messages.

ICMPv6 messages are subdivided into two classes: error messages and information messages.

Error Messages are divided into four categories:

1. Destination Unreachable
2. Time Exceeded
3. Packet Too Big
4. Parameter Problems

Information messages are divided into three groups:

1. Diagnostic messages
2. Neighbor Discovery messages
3. Messages for the management of multicast groups.

ICMPv6 Types and Codes

ICMPv6 has a number of messages that are identified by the "Type" field. Some of these types have assigned "Code" fields as well. The table below shows the different types of ICMP Types with their associated codes if

there are any.

Type codes 0 – 127 are error messages and type codes 128 – 255 are for information messages.

ICMPv6 Types and Codes

Type Number	Type Name	Code
0	Reserved	0 - no route to destination
		1 - communication with destination administratively prohibited
		2 - beyond scope of source address
		3 - address unreachable
		4 - port unreachable
		5 - source address failed ingress/egress policy
		6 - reject route to destination
		7 - Error in Source Routing Header
1	Destination Unreachable	
2	Packet Too Big	
3	Time Exceeded	0 - hop limit exceeded in transit
		1 - fragment reassembly time exceeded
4	Parameter Problem	0 - erroneous header field encountered
		1 - unrecognized Next Header type encountered
		2 - unrecognized IPv6 option encountered
100	Private Experimentation	
101	Private Experimentation	
102 - 126	Unassigned	
127	Reserved for expansion if ICMPv6 error messages	

Type Number	Type Name	Code
128	Echo Request	
129	Echo Replay	
130	Multicast Listener Query	
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	
134	Router Advertisement	
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
138	Router Renumbering	0 - Router Renumbering Command
		1 - Router Renumbering Result
		255 - Sequence Number Reset
139	ICMP Node Information Query	0 - The Data field contains an IPv6 address which is the Subject of this Query.
		1 - The Data field contains a name which is the Subject of this Query, or is empty, as in the case of a NOOP.
		2 - The Data field contains an IPv4 address which is the Subject of this Query.

Type Number	Type Name	Code
140	ICMP Node Information Response	0 - A successful reply. The Reply Data field may or may not be empty.
		1 - The Responder refuses to supply the answer. The Reply Data field will be empty.
		2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
141	Inverse Neighbor Discovery Solicitation Message	
142	Inverse Neighbor Discovery Advertisement Message	
143	Version 2 Multicast Listener Report	
144	Home Agent Address Discovery Request Message	
145	Home Agent Address Discovery Reply Message	
146	Mobile Prefix Solicitation	
147	Mobile Prefix Advertisement	
148	Certification Path Solicitation Message	
149	Certification Path Advertisement Message	

Type Number	Type Name	Code
150	ICMP messages utilized by experimental mobility protocols such as Seamoby	
151	Multicast Router Advertisement	
152	Multicast Router Solicitation	
153	Multicast Router Termination	
154	FMIPv6 Messages	
155	RPL Control Message	
156	ILNPv6 Locator Update Message	
157	Duplicate Address Request	
158	Duplicate Address Confirmation	
159 – 199	Unassigned	
200	Private experimentation	
201	Private experimentation	
255	Reserved for expansion of ICMPv6 informational messages	

IP

Internet Protocol (IP) is the primary part of the Network Layer of the OSI Model that is responsible for routing traffic across network boundaries. It is the protocol that is responsible for addressing. IPv4 is probable the version that most people are familiar with and it has been around since 1974. IPv6 is its current successor and due to a

shortage of available IPv4 addresses compared to the explosive increase in the number of devices that use IP addresses, IPv6 is rapidly increasing in use.

When IP is chosen as the protocol type the available option to further specify the protocol is the protocol number. This is used to narrow down which protocol within the Internet Protocol Suite and provide a more granular control.

Protocol Number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called "Protocol" to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the "Next Header" field.

Protocol Numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger

#	Protocol	Protocol's Full Name
16	CHAOS	Chaos
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol

#	Protocol	Protocol's Full Name
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6
61		any host internal protocol
62	CFTP	CFTP
63		any local network

#	Protocol	Protocol's Full Name
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol

#	Protocol	Protocol's Full Name
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol

#	Protocol	Protocol's Full Name
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield
131	PIPE	Private IP Encapsulation within IP
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	

#	Protocol	Protocol's Full Name
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching RFC 5237.

Protocol Number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called “Protocol” to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the “Next Header” field.

Protocol Numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway

#	Protocol	Protocol's Full Name
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger
16	CHAOS	Chaos
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol

#	Protocol	Protocol's Full Name
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header

#	Protocol	Protocol's Full Name
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonnet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6
61		any host internal protocol
62	CFTP	CFTP
63		any local network
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol

#	Protocol	Protocol's Full Name
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header

#	Protocol	Protocol's Full Name
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM

#	Protocol	Protocol's Full Name
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield
131	PIPE	Private IP Encapsulation within IP
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching RFC 5237.

Security policies

One of the foundations upon which a firewall works is the use of policies. These are what bring the other firewall objects and components together into an elegant mechanism for the governing of the traffic going through the network.

This Chapter includes information on the following topics:

- Firewall policies
- Security profiles
- SSL/SSH Inspection
- Identity Based Policies
- Device Identity Policies
- VPN Policies
- Interface Policies
- One-Arm IDS
- Local-In Policies
- Security Policy 0
- Deny Policies
- Accept Policies
- IPv6 Policies
- Fixed Port
- Endpoint Security
- Traffic Logging
- Quality of Service
- Policy Monitor

Firewall policies

The firewall policy is the axis around which most of the other features of the FortiGate firewall revolve. A large portion of the settings in the firewall at some point will end up relating to or being associated with the firewall policies and the traffic that they govern. Any traffic going through a FortiGate unit has to be associated with a policy. These policies are essentially discrete compartmentalized sets of instructions that control the traffic flow going through the firewall. These instructions control where the traffic goes, how it's processed, if it's processed and even whether or not it's allowed to pass through the FortiGate.

When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number). It also registers the incoming interface, the outgoing interface it will need to use and the time of day. Using this information the FortiGate firewall attempts to locate a security policy that matches the packet. If it finds a policy that matches the parameters it then looks at the action for that policy. If it is ACCEPT

the traffic is allowed to proceed to the next step. If the Action is DENY or a match cannot be found the traffic is not allowed to proceed.

The 2 basic actions at the initial connection are either ACCEPT or DENY:

- If the Action is ACCEPT, the policy action permits communication sessions. There may be other packet processing instructions, such as requiring authentication to use the policy. While you may not see it in the configuration there is the implied subset of the ACCEPT Action that include VPN policies, whether they be an IPsec VPN or SSL.
- If the Action is DENY, the policy action blocks communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped. A DENY security policy is needed when it is required to log the denied traffic, also called “violation traffic”.

The policy may contain a number of instructions for the FortiGate firewall in addition to the ACCEPT or DENY actions, some of which are optional. Instructions on how to process the traffic can also include such things as:

- Logging Traffic
- Authentication
- Network Address Translation or Port Address Translation
- Use Virtual IPs or IP Pools
- Caching
- Whether to use address or Identity based rules
- Whether to treat as regular traffic or VPN traffic
- What certificates to use
- Security profiles to apply
- Proxy Options
- Traffic Shaping

Firewall policy parameters

As mentioned before, for traffic to flow through the FortiGate firewall there must be a policy that matches its parameters:

Incoming Interface

This is the interface or interfaces that the traffic is first connection to the FortiGate unit by. The exception being traffic that the FortiGate generates itself. This is not limited to the physical Ethernet ports found on the device. The incoming interface can also be a logical or virtual interface such as a VPN tunnel, a Virtual WAN link or a wireless interface.

Outgoing Interface

After the firewall has processed the traffic it needs to leave a port to get to its destination and this will be the interface or interfaces that the traffic leaves by. This interface, like the **Incoming Interface** is not limited to only physical interfaces.

Source Address

The addresses that a policy can receive traffic from can be wide open or tightly controlled. For a public webserver that the world at large should be able to access, the best choice will be “all”. If the destination is a private webserver that only the branch offices of a company should be able to access or a list of internal computers that

are the only ones allowed to access an external resource then a group of preconfigured addresses is the better strategy.

Additional parameters under the Source Address, though they are not mandatory are:

- **Source User(s)**

This parameter is based on a user identity that can be from a number of authentication authorities. It will be an account or group that has been set up in advance that can be selected from the dropdown menu. The exception to this is the feature that allows the importing of LDAP Users. When the feature is used, a small wizard window will appear to guide the user through the setup. The caveat is that the LDAP server object in the **User and Device > Authentication > LDAP Servers** section has to be already configured to allow the use of this import feature.

- **Source Device Type**

This parameter is for narrowing down the traffic sending devices to those that the FortiGate is familiar with. Again the contents of this parameter need to be a preconfigured object and these are defined at **User and Device > Device > Device Definitions**. This parameter can limit the devices that can connect to this policy to those specific MAC addresses that are already known by the FortiGate and are approved for the policy.

Destination Address

In the same way that the source address may need to be limited, the destination address can be used as a traffic filter. When the traffic is destined for internal resources the specific address of the resource can be defined to better protect the other resources on the network. One of the specialized destination address options is to use a Virtual IP address. The destination address doesn't need to be internal you can define policies that are only for connecting to specific addresses on the Internet.

Schedule

The time frame that is applied to the policy. This can be something as simple as a time range that the sessions are allowed to start such as between 8:00 am and 5:00 pm. Something more complex like business hours that include a break for lunch and time of the session's initiation may need a schedule group because it will require multiple time ranges to make up the schedule.

Service

The service or service chosen here represent the TCP/IP suite port numbers that will most commonly be used to transport the named protocols or group of protocols. This will be a little different than Application Control which looks more closely at the packets to determine the actual protocol used to create them.

Without all six (possibly 8) of these things matching, the traffic will be declined. Each traffic flow requires a policy and the direction is important as well. Just because packets can go from point A to point B on port X does not mean that the traffic can flow from point B to point A on port X. A policy must be configured for each direction.

When designing a policy there is often reference to the traffic flow, but most communication is a two way connection so trying to determine the direction of the flow can be somewhat confusing. If traffic is HTTP web traffic the user sends a request to the web site, but most of the traffic flow will be coming from the web site to the user. Is the traffic flow considered to be from the user to the web site, the web site to the user or in both directions? For the purposes of determining the direction for a policy the important factor is the direction of the initiating communication. The user is sending a request to the web site so this is the initial communication and the web site is just responding to it so the traffic will be from the users network to the Internet.

A case where either side can initiate the communication like between two internal interfaces on the FortiGate unit would be a more likely situation to require a policy for each direction.

What is not expressly allowed is denied

One of the fundamental ideas that can be found in just about any firewall is the rule that anything that is not expressly allowed is by default denied. This is the foundation for any strategy of protecting your network. Right out of the box, once you have your FortiGate device connected into your network and hooked up with your ISP your network is protected. Nothing is getting out or in so it is not very convenient, but you don't have to worry that between the time you hooked it up and the point that you got all of the policies in place that someone could have gotten in and done something to your resources. The reason that this needs to be kept in mind when designing policies is because you cannot assume that any traffic will be allowed just because it makes sense to do so. If you want any kind of traffic to make it past the FortiGate firewall you need to create a policy that will allow that traffic. To maintain the protection of the network should also make sure that the any policy you create allows only the traffic you intend to go only to where you specifically want it to go and when you want it to go there.

Example

You have a web server on your network that is meant to provide a collaborative work environment web site for your employees and a partner company for a project over the course of the next 3 months.

It is theoretically possible to allow connections into your network to any device on that network for any service and at any time. The problem with this is that we might not want just anybody looking at those resources. Sadly, no matter how much it is wished otherwise, not everybody on the Internet can be trusted. Which means we now have to be very specific in our instructions as to what traffic to allow into the network. Each step that we take towards being more specific as to what we allow means that there is that much more that is not allowed and the level of protection of a resources is directly proportional to the amount of traffic that is not allowed. If somebody can't get at it they can't damage or steal it.

Limiting where the traffic is allowed to go to means that other computers on your network besides the web-server are protected.

- Limiting where the traffic is allowed to come from means that, if feasible, you can limit the systems that can access the web server to just employees or the partner company computers.
- Limiting the services to just web traffic means that a malicious person, even if they were connection from a computer at the partner organization could only use the features of web traffic to do anything malicious.
- Limiting the policy to the time span of the project would mean that even if the IT department forgot to remove the policy after the end of the project than no computer from the other company could be used to do anything malicious through the policy that allowed the traffic.

This is just a very basic example but it shows the underlying principles of how the idea that anything not expressly allowed is by default denied can be used to effectively protect your network.

Policy order

Another important factor in how firewall policies work is the concept of precedence of order or if you prefer a more recognizable term, "first come, first served".

It is highly likely that even after only a relatively small number of policies have been created that there will be some that overlap or are subsets of the parameters that the policies use to determine which policy should be matched against the incoming traffic. When this happens there has to be a method to determine which policy should be applied to the packet. The method which is used by most firewalls is based on the order of the sequence of the policies.

If all of the policies were placed in a sequential list the process to match up the packet would start at the top of the list and work its way down. It would compare information about the packet, specifically these points of information:

1. The interface the packet connected to the FortiGate firewall
2. The source address of the packet, and possibly the sending user and device.
3. The destination address of the packet
4. The interface the packet would need to use to get to the destination address based on the routing table
5. The port the packet is destined for
6. The time that the packet connected to the FortiGate

As soon as the a policy is reached that matches all of the applicable parameters, the instructions of that policy are applied and the search for any other matching policies is stopped. All subsequent policies are disregarded. Only 1 policy is applied to the packet.

If there is no matching policy among the policies that have been configured for traffic the packet finally drops down to what is always the last policy. It is an implicit policy. One of a few that are referred to by the term “policy0”. This policy denies everything.

The implicit policy is made up of the following settings:

- Incoming Interface: any
- Source Address: any
- Outgoing Interface: any
- Destination Address: any
- Action: DENY

The only setting that is editable in the implicit policy is the logging of violation traffic.

A logical best practice that comes from the knowledge of how this process works is to make sure that the more specific or specialized a policy is, the closer to the beginning of the sequence it should be. The more general a policy is the higher the likelihood that it could include in its range of parameters a more specifically targeted policy. The more specific a policy is, the higher the probability that there is a requirement for treating that traffic in a specific way.

Example

For security reasons there is no FTP traffic allowed out of a specific subnet so there is a policy that states that any traffic coming from that subnet is denied if the service is FTP, so the following policy was created:

Policy #1

Source Interface	Internal1
Source Address	192.168.1.0/24
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>

Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	deny

Now as these things usually go it turns out that there has to be an exception to the rule. There is one very secure computer on the subnet that is allowed to use FTP and once the content has been checked it can then be distributed to the other computer on the subnet. So a second firewall policy is created.

Policy #2

Source Interface	Internal1
Source Address	192.168.1.38/32
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>
Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	Allow

By default, a policy that has just been created will be placed last in the sequence so that it is less likely to interfere with existing policies before it can be moved to its intended position. If you look at Policy #2 you will notice that it is essentially the same as Policy #1 except for the Source Address and the Action. You will also notice that the Source Address of the Policy #2 is a subset of the Source address in policy #1. This means that if nothing further is done, Policy #2 will never see any traffic because the traffic will always be matched by Policy #1 and processed before it has a chance to reach the second policy in the sequence. For both policies to work as intended Policy #2 needs to be moved to before Policy #1 in the sequence.

Policy Identification

When looking at the policy listing it can appear as if the policies are identified by the sequence number in the far left column. The problem is that this number changes as the position of the policy in the sequence changes. The column that correctly identifies the policy, and the value sticks with the policy is the "ID" column. This column is not shown by default in the listing but can be added to the displayed columns by right clicking on the column heading bar and selecting it from the list of possible columns.

When looking in the configuration file the sequence is based upon the order of the policies as they are in the file just as they are in the list in the GUI. However, if you need to edit the policy in the CLI you must use the ID number.

UUID Support

Universally Unique Identifier (UUID) attributes have been added to policies to improve functionality when working with FortiManager or FortiAnalyzer units. If required, the UUID can be set manually through the CLI.

CLI Syntax:

```
config firewall {policy/policy6/policy46/policy64}
  edit 1
    set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
  next
end
```

Security profiles

Where security policies provide the instructions to the FortiGate unit for controlling what traffic is allowed through the device, the Security profiles provide the screening that filters the content coming and going on the network. Security profiles enable you to instruct the FortiGate unit about what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A security profile is a group of options and filters that you can apply to one or more firewall policies. Security profiles can be used by more than one security policy. You can configure sets of security profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same security profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Security profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure security profiles in the Security Profiles menu and applied when creating a security policy by selecting the security profile type.

There is a separate handbook for the topic of the Security Profiles, but because the Security Profiles are applied through the Firewall policies it makes sense to have at least a basic idea of what the security profile do and how they integrate into the FortiGate's firewall policies. The following is a listing and a brief description of what the security profiles offer by way of functionality and how they can be configured into the firewall policies.

AntiVirus

Antivirus is used as a catch all term to describe the technology for protection against the transmission of malicious computer code sometimes referred to as malware. As anyone who has listened to the media has heard that the Internet can be a dangerous place filled with malware of various flavours. Currently, the malware that is most common in the Internet, in descending order, is Trojan horses, viruses, worms, adware, back door exploits, spyware and other variations. In recent years, not only has the volume of malicious software become greater than would have been believed when it first appeared but the level of sophistication has risen as well.

The Antivirus Filter works by inspecting the traffic that is about to be transmitted through the FortiGate. To increase the efficiency of effort it only inspects the traffic being transmitted via the protocols that it has been configured to check. Before the data moves across the FortiGate firewall from one interface to another it is checked for attributes or signatures that have been known to be associated with malware. If malware is detected, it is removed.

Web Filtering

Malicious code is not the only thing to be wary of on the Internet. There is also the actual content. While the content will not damage or steal information from your computer there is still a number of reasons that would require protection from it.

In a setting where there are children or other sensitive people using the access provided by a connected computer there is a need to make sure that images or information that is not appropriate is not inadvertently displayed to them. Even if there is supervision, in the time it takes to recognize something that is inappropriate and then properly react can expose those we wish to protect. It is more efficient to make sure that the content cannot reach the screen in the first place.

In an organizational setting, there is still the expectation that organization will do what it can to prevent inappropriate content from getting onto the computer screens and thus provoking an Human Resources incident. There is also the potential loss of productivity that can take place if people have unfiltered access to the Internet. Some organizations prefer to limit the amount of distractions available to tempt their workers away from their duties.

The Web filter works primarily by looking at the destination location request for a HTTP(S) request made by the sending computer. If the URL is on a list that you have configured to list unwanted sites, the connection will be disallowed. If the site is part of a category of sites that you have configured to deny connections to the session will also be denied. You can also configure the content filter to check for specific key strings of data on the actual web site and if any of those strings of data appear the connection will not be allowed.

Application Control

Application Control is designed to allow you to determine what applications are operating on your network and to also filter the use of these applications as required. Application control is also for outgoing traffic to prevent the use of applications that are against an organization's policy from crossing the network gateway to other networks. An example of this would be the use of proxy servers to circumvent the restrictions put in place using the Web Filtering.

Intrusion Protection (IPS)

Intrusion Prevention System is almost self explanatory. In the same way that there is malware out on the Internet that the network needs to be protected from there are also people out there that take a more targeted approach

to malicious cyber activity. No operating system is perfect and new vulnerabilities are being discovered all of the time. An intrusion prevention system is designed to look for activity or behavior that is consistent with attacks against your network. When attack like behavior is detected it can either be dropped or just monitored depending on the approach that you would like to take.

As new vulnerabilities are discovered they can be added to the IPS database so that the protection is current.

Email Filtering

Spam or unsolicited bulk email is said to account for approximately 90% of the email traffic on the Internet. Sorting through it is both time consuming and frustrating. By putting an email filter on policies that handle email traffic, the amount of spam that users have to deal with can be greatly reduced.

Data Leak Prevention (DLP)

Data Leak Prevention is used to prevent sensitive information from leaving your network. When people think of security in the cyber-world one of the most common images is that of a hacker penetrating your network and making off with your sensitive information, but the other way that you can lose sensitive data is if someone already on the inside of your network sends it out. This does not have to be an act of industrial espionage. It can just be a case of not knowing the policies of the organization or a lack of knowledge of security or laws concerning privacy.

For instance, a company may have a policy that they will not reveal anyone's Social Security number, but an employee emails a number of documents to another company that included a lengthy document that has a Social Security number buried deep within it. There is not malicious intent but if the information got out there could be repercussions.

If an organization has any information in a digital format that it cannot afford for financial or legal reasons, to leave its network, it makes sense to have Data Leak Prevention in place as an additional layer of protection.

VoIP

Voice over IP is essentially the protocols for transmitting voice or other multimedia communications over Internet Protocol networks such as the Internet. The Security Profiles VoIP options apply the SIP Application Level Gateway (ALG) to support SIP through the FortiGate unit. The SIP ALG can also be used to protect networks from SIP-based attacks.

ICAP

Internet Content Adaptation Protocol (ICAP) off loads HTTP traffic to another location for specialized processing. The purpose of this module when triggered is to send the incoming HTTP traffic over to a remote server to be processed thus taking some of the strain off of the resources of the FortiGate unit. The reasons for the specialized process could be anything from more sophisticated Antivirus to manipulation of the HTTP headers and URLs.

EndPoint Control

EndPoint Control makes sure that certain standards are kept. When a computer on the Internet becomes connected to the FortiGate unit by VPN that computer is now part of the same network and therefore needs to be

subject to the same levels of protection, not only to protect the computer but the network. In the EndPoint Control section you can set the minimum standards for things like AntiVirus software and VPN software.

Proxy Option Components

Any time a security profile that requires the use of a proxy is enabled the Proxy Options field will be displayed. Certain inspections defined in security profiles require that the traffic be held in proxy while the inspection is carried out and so the Proxy Options are there to define the parameters of how the traffic will be processed and to what level the traffic will be processed. In the same way that there can be multiple security profiles of a single type there can also be a number of unique Proxy Option profiles so that as the requirements for a policy differ from one policy to the next you can also configure a different Proxy Option profile for each individual policy or you can use one profile repeatedly.

The Proxy Options refer to the handling of the following protocols:

- HTTP
- SMTP
- POP3
- IMAP
- FTP
- NNTP
- MAPI
- DNS
- IM

The configuration for each of these protocols is handled separately.

The use of different proxy profiles and profile options

Just like other components of the FortiGate, there is the option for different Proxy Option profiles so that you can be very granular in your control of the workings of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

Oversized File Log

This setting is for those that would like to log the occurrence of oversized files being processed. It does not change how they are processed it only enables the FortiGate unit to log that they were either blocked or allowed through. A common practice is to allow larger files through without antivirus processing. This allows you to get an idea of how often this happens and decide on whether or not to alter the settings relating to the treatment of oversized files.

The setting of the threshold for what is considered to be an oversized file is located in the Oversized File / Email Threshold that is found in some of the protocol options for the Proxy Options.

Protocol Port Mapping

While each of the protocols listed has a default TCP port that is commonly used, the level of granularity of control on the FortiGate firewall allows that the port used by the protocols can be individually modified in each separate

Profile. It can also be set to inspect any port with flowing traffic for that particular protocol. The headers of the packets will indicate which protocol generated the packet. To optimize the resources of the unit the mapping and inspection of protocols can be enabled or disabled depending on your requirements.

Comfort Clients

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit begins scanning the file. During the buffering and scanning procedure, the user must wait. After the scan is completed, if no infection is found, the file is sent to the next step in the process flow. If the file is a large one this part of the process can take some time. In some cases enough time that some users may get impatient and cancel the download.

The comfort client feature mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete so as to let the user know that processing is taking place and that there hasn't been a failure in the transmission. This slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned without any indication of viruses the transfer will proceed at full speed.

If there is evidence of an infection the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file. If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Oversized File/Email Threshold

This is another feature that is related to antivirus scanning. The FortiGate unit has a finite amount of resources that can be used to buffer and scan a file. If a large file such as an ISO image or video file was to be downloaded this could not only overwhelm the memory of the FortiGate, especially if there were other large files being downloaded at the same time, but could exceed it as well. For this reason, how to treat large files needs to be addressed.

A threshold is assigned to determine what should be considered an oversize file or email. This can be set at any size from 1 MB to 50 MB. Any file or email over this threshold will not be processed by the Antivirus Security Profiles. Once a file is determined to be oversized it must be then determined whether to allow it or to block it.

These settings are not a technical decision but a policy one that will depend on your comfort level with letting files into your network. As there often is, there is a compromise between convenience or ease of use and security. If you want to go for a high peace of mind level you can configure the firewall to block oversized files and thus no files would be coming into the network that have not been scanned. If you are looking for optimizing the memory of the FortiGate unit and making sure that everybody is getting the files they want, you can lower the threshold and allow files that are over the threshold.

It should be noted that in terms of probability that malware is more likely to be found in smaller files than in larger files. A number of administrators take this into account when they lower the default threshold so as to lessen the impact on memory if they see the FortiGate unit going into conserve mode on a regular basis.

Chunked Bypass

The HTTP section allows the enabling of “Chunked Bypass”. This refers to the mechanism in version 1.1 of HTTP that allows a web server to start sending chunks of dynamically generated output in response to a request before actually knowing the actual size of the content. Where dynamically generated content is concerned this means that there is a faster initial response to HTTP requests. From a security stand point it means that the content will not be held in the proxy as an entire file before proceeding.

Allow Fragmented Messages

The specifications of RFC 2046 allow for the breaking up of emails and sending the fragments in parallel to be rebuilt and read at the other end by the mail server. It was originally designed to increase the performance over slower connections where larger email messages were involved. It will depend on your mail configuration if this is even possible for your network but outside of Microsoft Outlook and Outlook Express, not many email clients are set up to break up messages like this. The drawback of allowing this feature is that if malware is broken up between multiple fragments of the message the risk is run that it will not be detected by some antivirus configurations because the code may not all be present at the same time to identify.

Append Email Signature

The Append Email Signature is used when an organization would like to ensure that over and above our in this case underneath the existing personal signatures of the sender, all of the emails going out of their network have the appropriate “boilerplate”, for lack of a better term. These appended emails do not replace existing signatures. They are as the feature states, appended to the email.

Examples could include things like:

- Without prior approval the email should not be forwarded.
- Please be environmentally friendly and don't print out emails
- For questions regarding the purchasing of our products please call...

It can be anything that the organization would like as long as it is in text format. The use of this feature usually works best in an environment where there is some standardization of what goes into the personal signatures of the senders so that there is no duplication or contradiction of information in the signatures.

Viewing Firewall Policies

When you first go into the Policy window, found by going to Policy > Policy > Policy, you will see a table with a menu bar across the top. The menu bar will have the following items:

At the top left:

- Create New (with a “+” sign on the left and a downward pointing triangle on the right)
- Clone
- Delete
- Column Settings
- Filter Settings

At the top right:

- Section View
- Global View

The items at the top right with their radio buttons represent the 2 potential views that the policies can be displayed in.

The Global View shows all of the policies in the order of their sequence. With the default settings you will be able to see the sequence number in a column close to the left side of the table.

The Section view is similar to the Global View except that as the name implies it is divided into sections. By default the sections are based on the paths between the interfaces. These can be referred to as “interface pairings”. For instance, all of the policies referencing traffic from WAN1 to DMZ will be in one section. The policies referencing traffic from DMZ to WAN1 will be in another section.

The sections are collapsible so that you only need to look at the sections with policies you are interested in. It is possible to add customized subsections within the default sections of interface pairings. This would be useful in a situation where you have a lot of policies and would like to further compartmentalize them by common attributes so that things are easier to find.

The default column headings are:

- [Check box icon]
- Seq.#
- Source
- Destination
- Authentication
- Schedule
- Service
- Action
- Log

The columns that are shown are configurable. All but the first 2 can be removed or their position changed. There are also a number of other columns that display information about the policies that can be added. One of the more useful ones that can be added is the ID column. The reason for adding this one is that policies are referenced by their ID number for simplicity and ease of administration. If you are looking in the CLI you will see that the only designation for a policy is its number and if you wish to change the order of a policy you will be asked to move it before or after another policy by referencing its number.

How “Any” policy can remove the Section View

The FortiGate unit will automatically change the view on the policy list page to Global View whenever a policy containing “any” in the Source interface/zone or Destination interface/zone is created. If the Section View is greyed out it is likely that one or more of the policies has “any” as a Source or Destination interface.

With the use of the “any” the policy should go into multiple sections because it could effectively be any of a number of interface pairings. As mentioned, policies are sectioned by using the interface pairings (for example, port1 -> port2) and each section has its own specific policy order. The order in which a policy is checked for matching criteria to a packet’s information is based solely on the position of the policy within its section or within the entire list of policies as a whole but if the policy is in multiple sections at the same time there is no mechanism for placing the policy in a proper order within all of those sections at the same time because it is a manual process and there is no parameter to compare the precedence of one section or policy over the other. Thus a conflict is

created. In order to resolve the conflict the FortiGate firewall removes that aspect of the sections so that there is no need to compare and find precedence between the sections and it therefore has only the Global View to work with.

Security policy configuration extensions

When first creating the policy the configuration form will ask for a choice between the policy types of Firewall or VPN, Firewall being the default. Choosing whether or not to leave the selection as Firewall is straight forward. If the policy is not a policy based VPN policy then it is a Firewall policy type.

There are essentially 2 types of VPN connections, Interface Based and Policy Based. In an Interface Based VPN tunnel a logical interface is created that can be seen as an interface by the policies in the same way that any of the physical interfaces can be seen. Therefore to govern the traffic a regular policy will work. The policy based VPN tunnels work slightly different and therefore need a slightly different policy configuration. For a more detail explanation of the difference between the types of VPN tunnels refer to the VPN documentation found in the VPN handbooks or in the VPN section of the Complete Administration Guide.

Once either the Firewall or the VPN type has been chosen there is then a choice between one of subtypes for each of the Policy types. For the Firewall type of policy the subtypes are:

- Address
- User Identity
- Device Identity

The Address subtype refers to policies where access through the FortiGate firewall is dependant on the source location of the addresses of the devices involved in the traffic matched to the policy.

The User Identity subtype refers to policies where access through the FortiGate firewall is dependant on the users credentials or Identity.

The Device Identity subtype refers to policy where access through the FortiGate firewall is dependant on the specific device being used based on the MAC address of the device or belonging to a group of devices that are based on device types or belonging to custom made groups.

For the VPN type the subtypes are:

- IPsec
- SSL-VPN

As expected the two subtypes are the two different types of VPN tunnels that the FortiGate firewall supports in a policy based configuration.

SSL/SSH Inspection

While the profile configuration for this is found in the Security Profiles section it is enabled in the firewall policy along with the security profiles. This sort of analysis is some times referred to as deep scanning.

Deep Inspection works along the following lines. If your FortiGate unit has the correct chipset it will be able to scan SSL encrypted traffic in the same way that regular traffic can be scanned. The FortiGate firewall will essentially receive the traffic on behalf of the client and open up the encrypted traffic. Once it is finished it re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. By enabling this feature, it allows the FortiGate firewall to filter on traffic that is using the SSL encrypted protocol.

The encrypted protocols that can be inspected are:

- HTTPS
- SMTPS
- POP3S
- IMAPS
- FTPS

Before the invention of SSL inspection, scanning regular web traffic can be circumvented by using the prefix `https://` instead of `http://` in the URL. SSL inspection prevents this circumvention. However, because when the encrypted traffic is decrypted it has to be re-encrypted with the FortiGate's certificate rather than the original certificate it can cause errors because the name on the certificate does not match the name on the web site.

At one point deep inspection was something that was either turned on or off. Now individual deep inspection profiles can be created depending on the requirements of the policy. Depending on the Inspection Profile, you can:

- Configure which CA certificate will be used to decrypt the SSL encrypted traffic.
- Configure which SSL protocols will be inspected.
- Configure which ports will be associated with which SSL protocols for the purpose of inspection.
- Configure which websites will be exempt from SSL inspection
- Configure whether or not to allow invalid SSL certificates.
- Configure whether or not SSH traffic will be inspected.

HTTP Strict Transport Security (HSTS) Protocol

HSTS is a protocol used by Google and other web browsers to prevent man-in-the-middle attacks.

When performing deep inspection, the FortiGate intercepts the https traffic and would send its own self-signed CA certificate to the browser. If the browser is configured to use HSTS connections, it would refuse the FortiGate CA certificate since it is not on the trusted list for Google servers.

To keep the CA certificate from being refused, the HSTS settings should be cleared from the browser. Instructions for this vary between browsers.

Inspection Exemption

When you are using a browser to visit SSL encrypted sites and we are using a certificate that does not match the certificate of the site, we are presented with a warning message and the option of continuing, using the untrusted certificate, or terminating the session. However, there are a number of applications that use SSL encrypted traffic. If the application detects SSL traffic that wasn't signed with a certificate that it trusts it will not allow the traffic. The applications do not give the option to manually indicate that we trust the certificate or the site.

If the option is available, the customer may choose to import needed SSL certificates into Local Certificates and configure a policy for communication for that application.

The assist in preventing loss of access to these site but still enabling the SSL inspection of the rest of the internet traffic, a method of exempting either Website categories or specific sites has been developed. To exempt a large group of sites the profile can be configure to exempt FortiGuard Categories. There are 3 of these categories preselected due to the high likelihood of issues with associated applications with the type of websites included in these categories.

- Heath and Wellness
- Personal Privacy
- Finance and Banking

Other more specific websites can be added to the exemption list by creating addresses for them at **Policy & Objects > Objects > Addresses**. The adding of addresses is done by selection from a drop down menu. There is an option at the bottom of the list to create a new address, but otherwise only preconfigured addresses that are configured to be on the "Any" interface will be available for selection.

Examples of sites that you may want to configure for exemption so that there will be no interference due to certificate issues:

Apple

- *.appstore.com
- *.apple.com
- *.itunes.apple.com
- *.icloud.com
- swscan.apple.com

Dropbox

- *.dropbox.com

Skype

- *.messenger.live.com

Windows Updates

- update.microsoft.com

Allow Invalid SSL Certificate

This setting was something that used to be part of the **Proxy Options**, but now that SSL inspection has it's own configuration setting it is configured with those. It might seem like a straight forward decision that the allowing of invalid SSL certificates must be bad and therefore should not be allowed, but there can be some reasons that should be considered. The issues at hand are the reasons to use a SSL certificate and the reasons that a certificate will be considered invalid.

At a purely technical level, a properly formed certificate will encrypt the data so that it can only be read by the intended parties and not be read by anyone sniffing traffic on the network. For this reason, people will often use self-signed certificates. These self signed certificates are free and will encrypt the data just as well as those purchased from any of the big vendors of certificates, but if they are not listed as an approved Certificate Authority (CA) the certificates will be considered invalid.

On the other hand, one of the services the vendors provide is verification of identity of those that purchase their certificates. This means that if you see a valid certificate from a site that identified itself as being from "valid-company.com" that you can be reasonably sure that the site does belong to that company and not a false site masquerading as being part of that company.

Identity Based Policies

Identity based policies are ones in which there is the additional component of either an account identity or device identity. The inclusion of one or both of these components adds an extra dimension of complexity to working with these policies in the context of the other policies so while the extra security and granularity of control are beneficial, extra care must be taken when configuring the policies themselves and how they are positioned in the policy sequence. The actual configuration of these identities are explained in detail in the Authentication Handbook.

Identity-based security policies are usually configured for IPsec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

Identity-based policy positioning

In non-identity based policies, if non of the 6 mandatory policy parameters matches the header of the traffic packets the parameters are compared against the next policy in sequence. Because those parameters are mandatory there is always a value to test against and whether or not the policy applies is certain. The fact that the identity parameters are not required makes knowing whether or not the correct policy will be applied less obvious.

Originally, the identity aspect of a policy was an entire sub-policy checking sequence within each policy, including its own 0 policy at the end of the sequence. If all of the other parameters match the policy would then compare the traffic's identity with the list of identity groups in the policy starting at the beginning of the sequence and going through them until an identity was found that matched and then the rules for that identity group would be applied. If the traffic's identity did not match any of those listed in the policy it go to the last identity in the policy would be everyone and the Action would be deny.

The identity aspects of policies have now been incorporated in a single flat configuration that makes them a fundamental part of the policy rather than something that is added to the policy. This is simpler and allows for more complex combinations of address identification, user authentication and device determination that were not possible with previous policy configurations. Both user groups and device groups can be part of the same policy. Because the identity aspects are optional, more flexibility in creating policies that use authentication is possible.

Identity fall through rules

The fall through rules for policies in 5.2 have changed so that they are more in keeping with the practices of other vendors. This makes it easier for users used to other firewalls to configure the policies and it also makes it simpler to convert the policies of other firewalls to be used on a FortiGate firewall.

Previously, if traffic reached an identity policy and the user or device was not a member of one of the groups specified it would fall through to the implicit deny all policy. This meant that any traffic that reached that policy would have to be authenticated and a member of one of the listed groups. If the 6 required parameters matched, the traffic would not be getting past this policy.

The approach is now to treat the the identity parameters, if they exist, the same as the other parameters, in that if they do not match any listed in the policy, the traffic drops down to the next policy.

Example:

There are three policies where all the parameters are the same except:

- Policy # 1 - Source User Group A is assigned profile A
- Policy # 2 - Source User Group B s assigned profile B

- Policy # 3 - Source User(s) and Source Device Type are empty

Traffic that matches all of the required parameters will be processed as follows:

- Traffic authenticated as being from User Group A will be processed by Policy # 1.
- Traffic authenticated as being from User Group B will be processed by Policy # 2.
- Traffic with no authenticated users will be processed by Policy # 3.
- Traffic authenticated as being from User Group C will be processed by Policy # 3.

In the methodology before FortiOS 5.2, traffic authenticated as being User Group B, User Group C or no authenticated user at all would have been stopped at Policy # 1.

The CLI command “fall-through-unauthenticated” that was added in 5.0.1 attempted to allow a process similar to this, but only applied to unauthenticated traffic and not authenticated traffic that didn’t match the list of groups is the the sub-policy. The current methodology is not subject to the same limitation and alleviates the need for the function of this command so the command has been removed from the CLI.

Implicit Protocols

In previous versions of the firmware, the protocols that were used to authenticate such as HTTP, HTTPS, FTP, and Telnet, were supported on the policy whether or not they were included in the supported services. In 5.2, the protocol needed to authenticate needs to be included in the list of allowed services in order the the authentication to take place.

For example, if you have a VIP coming into your network that is for connecting to some security webcams located in your data center that use custom services or ports to connect to, if you are using an identity policy you would also have to include HTTP or HTTPS in the services list in order to actually authenticate.

Another formerly implicit protocol that is not supported automatically in 5.2 is port 53 (DNS). If you are limiting the services of a protocol to web based protocols such as HTTP or HTTPS don't forget to add DNS so that the domain names can be resolved.

When upgrading the firmware from version 5.0.x to 5.2.x, a policy with either an identity or device sub-policy will automatically convert from a single policy with sub-policies to a separate policy for each identity based sub-policy.

VPN Policies

At one point, if you wanted to have secure digital communications between 2 points a private network would be created. This network would only allow the people that were intended to get the communications on it. This is very straightforward if the 2 points are in the same room or even in the same building. It can all be done physically. If you are supposed to be on the secure network

VPNs are an answer to one of today's biggest concerns, how to make digital communications secure between to points that must communicate over the Internet which anybody can have access to

IPsec Policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate firewall interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate firewall interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.

For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

DoS Protection

Denial of Service (DoS) policies are primarily used to apply DoS anomaly checks to network traffic based on the FortiGate interface it is entering as well as the source and destination addresses. DoS checks are a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and behavior. A common example of anomalous traffic is the denial of service attack. A denial of service occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system, so that legitimate users can no longer use it.

DoS policies are similar to firewall policies except that instead of defining the way traffic is allowed to flow, they keep track of certain traffic patterns and attributes and will stop traffic displaying those attributes. Further, DoS policies affect only incoming traffic on a single interface. You can further limit a DoS policy by source address, destination address, and service.

DoS configurations have been changed a couple of times in the past. In FortiOS 4.0, DoS protection is moved to the interface policy, so when it is enabled, it is the first thing checked when a packet enters FortiGate. Because of this early detection, DoS policies are a very efficient defence that uses few resources. Denial of service attacks, for example, are detected and its packets dropped before requiring security policy look-ups, antivirus scans, and other protective but resource-intensive operations.

A DoS policy examines network traffic arriving at an interface for anomalous patterns usually indicating an attack. This does not mean that all anomalies experience by the firewall are the result of an intentional attack.

Because an improperly configured DoS anomaly check can interfere with network traffic, no DoS checks are preconfigured on a factory default FortiGate unit. You must create your own before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.

To create a Denial of Service policy determine if it needs to be an IPv4 or IPv6 policy, then goto:

Policy & Objects > Policy > DoS Policy for IPv4.

Policy & Objects > Policy > IPv6 DoS Policy for IPv6.

The **Enable SSH Deep Scan** feature is enabled by default when creating a new SSL/SSH Inspection profile. There are situations where this feature can cause issues so be sure that you would like it enabled before applying it.

Settings used in configuring DoS

Incoming Interface

The interface to which this security policy applies. It will be the that the traffic is coming into the firewall on.

Source Address

This will be the address that the traffic is coming from and must be a address listed in the Address section of the Firewall Objects. This can include the predefined “all” address which covers any address coming in on any interface. Multiple addresses or address groups can be chosen

Destination Address

This will be the address that the traffic is addressed to. In this case it must be an address that is associated with the firewall itself. For instance it could be one of the interface address of the firewall, a secondary IP address or the interface address assigned to a Virtual IP address. Just like with the Source Address this address must be already configured before being used in the DoS policy. Multiple addresses, virtual IPs or virtual IP groups can be chosen.

Service

While the Service field allows for the use of the ALL service some administrators prefer to optimize the resources of the firewall and only check on the services that will be answered on an interface. Multiple services or service groups can be chosen.

Anomalies

The anomalies can not be configured by the user. They are predefined sensors set up for specific patterns of anomalous traffic

The anomalies that have been predefined for use in the DoS Policies are:

Anomaly Name	Description	Recommended Threshold
tcp_syn_flood	If the SYN packet rate of new TCP connections, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
tcp_port_scan	If the SYN packet rate of new TCP connections, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second.

Anomaly Name	Description	Recommended Threshold
tcp_src_session	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
tcp_dst_session	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_flood	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_scan	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_src_session	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_dst_session	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
icmp_flood	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	250 packets per second.
icmp_sweep	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed.	100 packets per second.
icmp_src_session	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.	300 concurrent sessions
icmp_dst_session	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.	3000 concurrent sessions
ip_src_session	If the number of concurrent IP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
ip_dst_session	If the number of concurrent IP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.

Anomaly Name	Description	Recommended Threshold
sctp_flood	If the number of SCTP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second
sctp_scan	If the number of SCTP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second
sctp_src_session	If the number of concurrent SCTP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions
sctp_dst_session	If the number of concurrent SCTP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions

Status

The status field is enabled to enable the sensor for the associated anomaly. In terms of actions performed there is no difference between disabling a sensor and having the action as "Pass" but by disabling sensors that are not being used for blocking or logging you can save some resources of the firewall that can be better used elsewhere.

Logging

Regardless of whether the traffic is blocked or passed through the anomalous traffic will be logged.

Pass

Allows the anomalous traffic to pass through unimpeded.

Block

For Thresholds based on the number of concurrent sessions blocking the anomaly will not allow more than the number of concurrent sessions set as the threshold.

For rate based thresholds where the threshold is measured in packets per second, the Action setting "Block" prevents the overwhelming of the firewall by anomalous traffic in one of 2 ways. Setting which of those 2 ways will be issued is determined in the CLI.

- continuous - blocks any packets that match the anomaly criteria once the threshold has been reached
- periodical - allows matching anomalous traffic up to the rate set by the threshold.



If the period for a particular anomaly is 60 seconds, such as those where the threshold is measured in concurrent sessions, after the 60 second timer has expired, the number of allowed packets that match the anomaly criteria is reset to zero. This means that if you allow 10 sessions through before blocking, after the 60 seconds is up, another 10 will be allowed. The attrition of sessions from expiration should keep the allowed sessions from reaching the maximum.

To set the type of block action for the rate based anomaly sensors:

```
config ips global
    set anomaly-mode continuous
    set anomaly-mode periodical
end
```

Threshold

The threshold can be either in terms of concurrent session or in packets per second depending on which sensor is being referred to.

One-Arm IDS

Interface-based policy only defines what and how IPS functions are applied to the packets transmitted by the interface. It works no matter if the port is used in a forwarding path or used as an One-Arm device.

To enable One-Arm IDS, the user should first enable sniff-mode on the interface,

```
config system interface
    edit port2
        set ips-sniffer-mode enable
    next
end
```

Once sniff-mode is turned on, both incoming and outgoing packets will be dropped after IPS inspections. The port can be connected to a hub or a switch's SPAN port. Any packet picked up by the interface will still follow the interface policy so different IPS and DoS anomaly checks can be applied.

IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create an normal IPS sensor and assign it to the IPv6 interface policy.

```
config firewall interface-policy6
    edit 1
        set interface "port1"
        set srcaddr6 "all"
        set dstaddr6 "all"
        set service6 "ANY"
        set ips-sensor-status enable
        set ips-sensor "all_default"
    next
end
```

Traffic Destined to the FortiGate unit

IPS enabled in firewall policies can only inspect the traffic pass through FortiGate unit, not the traffic destined to FortiGate unit. Enabling IPS in interface-policy allows IPS to pick up any packet on the interface so it is able to inspect attacks targeting FGT.

Dropped, Flooded, Broadcast, Multicast and L2 packets

In many evaluation or certification tests, FortiGate firewall is often required to log any packets dropped by the firewall. In most of cases, these packets are of invalid headers so firewall just drops them silently. It is natural to forward all these packets to IPS first so FortiGate firewall is able to generate logs for invalid packets.

Flooded, broadcast and multicast traffics do not reach any of services in the forwarding path. They can be inspected by the interface policy as long as they match the addresses defined. Potentially, L2 packets can also be sent to IPS for inspection through interface-policy, but it is not enabled in FortiOS 4.0.

GUI and CLI

Now in FortiGate, there are two places that IPS can be enabled, in a firewall policy and in an interface policy. In the firewall policy implementation, IPS sensor can be configured in both CLI and GUI. When adding an IPS sensor to an interface policy it must be done through the CLI. There is no GUI input window for the “Interface Policy”. There is however, a DoS Policy section in the GUI.

Local-In Policies

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
  edit <policy_number>
    set intf <source_interface>
    set srcaddr <source_address>
    set dstaddr <destination_address>
    set action {accept | deny}
    set service <service name>
    set schedule <schedule_name>
  end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12, represented by the address object mgmt-comp1, using SSH on port 3 (192.168.21.77 represented by the address object FG-port3) using the Weekend schedule which defines the time the of access.

```
config firewall local-in-policy
edit <1>
set intf port3
set srcaddr mgmt-comp1
set dstaddr FG-port3
set action accept
set service SSH
set schedule Weekend
end
```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```
config firewall local-in-policy
edit <policy_number>
set status disable
end
```

Use the same commands with a status of enable to use the policy again.

Local-in policies are also supported for IPv6 by entering the command

```
config firewall local-in-policy6.
```

Security Policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate firewall logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic subtype=violation pri=warning
vd=root SN=179089 duration=0 user=N/A group=N/A rule=0 policyid=0 proto=17
service=137/udp app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A dst_int="Internal" sent=0 rcvd=0
src_port=137 dst_port=137 vpn=N/A tran_ip=0.0.0.0 tran_port=0
```

Deny Policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.

Accept Policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

Fixed Port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable fixedport when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
  edit <policy-id>
    ...
    set fixedport enable
    ...
  end
```

However, enabling fixedport means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the

Endpoint Security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the Security Profiles chapter in the FortiOS Handbook.

Traffic Logging

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Depending on what the FortiGate unit has in the way of resources, there may be advantages in optimizing the amount of logging taking places. This is why in each policy you are given 3 options for the logging:

- **No Log** - Does not record any log messages about traffic accepted by this policy.
- **Log Security Events** - records only log messages relating to security events caused by traffic accepted by this policy.
- **Log all Sessions** - records all log messages relating to all of the traffic accepted by this policy.

Depending on the the model, if the Log all Sessions option is selected there may be 2 additional options. These options are normally available in the GUI on the higher end models such as the FortiGate 600C or larger.

- **Generate Logs when Session Starts**
- **Capture Packets**

You can also use the CLI to enter the following command to write a log message when a session starts:

```
config firewall policy
  edit <policy-index>
    set logtraffic-start
  end
```

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13
05:23:47
log_id=4
type=traffic
subtype=other
pri=notice
vd=root
status="start"
src="10.41.101.20"
srcname="10.41.101.20"
src_port=58115
dst="172.20.120.100"
dstname="172.20.120.100"
dst_country="N/A"
dst_port=137
tran_ip="N/A"
tran_port=0
tran_sip="10.31.101.41"
```

```

tran_sport=58115
service="137/udp"
proto=17
app_type="N/A"
duration=0
rule=1
policyid=1
sent=0
rcvd=0
shaper_drop_sent=0
shaper_drop_rcvd=0
perip_drop=0
src_int="internal"
dst_int="wan1"
SN=97404 app="N/A"
app_cat="N/A"
carrier_ep="N/A"

```

If you want to know more about logging, see the Logging and Reporting chapter in the FortiOS Handbook. If you want to know more about traffic log messages, see the FortiGate Log Message Reference.

Quality of Service

The Quality of Service (QoS) feature allows the management of the level of service and preference given to the various types and sources of traffic going through the firewall so that the traffic that is important to the services and functions connecting through the firewall gets the treatment required to ensure the level of quality that is required.

QoS uses the following techniques:

Traffic policing	Packets are dropped that do not conform to bandwidth limitations
Traffic Shaping	Assigning minimum levels of bandwidth to be allocated to specific traffic flows to guarantee levels of servers or assigning maximum levels of bandwidth to be allocated to specific traffic flows so that they do not impede other flows of traffic.

This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.

Queuing

Assigning differing levels priority to different traffic flows so that traffic flows that are adversely effected by latency are prevented from being effected by traffic flows that are not subject to the effects of latency. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

An example of where you would want to use something like this is if you had competing traffic flows of Voice over IP traffic and email traffic. The VoIP traffic is highly susceptible to latency issues. If you have a delay of a few seconds it is quickly noticeable when it is occurring. Email on the other hand can have a time delay of much longer and it is highly unlikely that it will be noticed at all.

By default, the priority given to any traffic is high, so if you want to give one type of traffic priority over all other traffic you will need to lower the priority of all of the other traffic.

Policy Monitor

Once policies have been configured and enabled it is useful to be able to monitor them. To get an overview about what sort of traffic the policies are processing go to Policy > Monitor > Policy Monitor.

The window is separated into two panes.

Upper Pane

The upper pane displays a horizontal bar graph comparing the **Top Policy Usage** based on one of the following criteria:

- Active Sessions
- Bytes
- Packets

The criteria that the displayed graph is based on can be selected from the drop down menu in the upper right corner of the pane. The field name is **Report By**.

The bars of the graph are interactive to an extent and can be used to drill down for more specific information. If you hover the cursor over the bar of the graph a small popup box will appear displaying more detailed information. If the bar of the graph is selected an entirely new window will be displayed using a vertical bar graph to divide the data that made up the first graph by IP address.

For example if the first graph was reporting usage by active sessions it would include a bar for each of the top policies with a number at the end showing how many sessions were currently going through that policy. If one of the bars of the graph was then selected the new bar graph would show the traffic of that policy separated by either **Source Address**, **Destination Address** or **Destination Port**. As in the other window, the selection for the reported criteria is in the upper right corner of the pane. If the parameter was by source address there would be a bar for each of the IP addresses sending a session through the policy and the end of the bar would show how many sessions.

To go back to the previous window of information in the graphs select the Return link in the upper left of the pane.

Lower Pane

The lower pane contains a spreadsheet of the information that the bar graph will derive their information from. The column headings will include:

- Policy ID
- Source Interface/Zone
- Destination Interface/Zone
- Action
- Active Sessions

- Bytes
- Packets

Network defense

This section describes in general terms the means by which attackers can attempt to compromise your network and steps you can take to protect it. The goal of an attack can be as complex as gaining access to your network and the privileged information it contains, or as simple as preventing customers from accessing your web server. Even allowing a virus onto your network can cause damage, so you need to protect against viruses and malware even if they are not specifically targeted at your network.

The following topics are included in this section:

- Monitoring
- Blocking external probes
- Defending against DoS attacks

Monitoring

Monitoring, in the form of logging, alert email, and SNMP, does not directly protect your network. But monitoring allows you to review the progress of an attack, whether afterwards or while in progress. How the attack unfolds may reveal weaknesses in your preparations. The packet archive and sniffer policy logs can reveal more details about the attack. Depending on the detail in your logs, you may be able to determine the attackers location and identity.

While log information is valuable, you must balance the log information with the resources required to collect and store it.

Blocking external probes

Protection against attacks is important, but attackers often use vulnerabilities and network tools to gather information about your network to plan an attack. It is often easier to prevent an attacker from learning important details about your network than to defend against an attack designed to exploit your particular network.

Attacks are often tailored to the hardware or operating system of the target, so reconnaissance is often the first step. The IP addresses of the hosts, the open ports, and the operating systems the hosts are running is invaluable information to an attacker. Probing your network can be as simple as an attacker performing an address sweep or port scan to a more involved operation like sending TCP packets with invalid combinations of flags to see how your firewall reacts.

Address sweeps

An address sweep is a basic network scanning technique to determine which addresses in an address range have active hosts. A typical address sweep involves sending an ICMP ECHO request (a ping) to each address in an address range to attempt to get a response. A response signifies that there is a host at this address that responded to the ping. It then becomes a target for more detailed and potentially invasive attacks.

Address sweeps do not always reveal all the hosts in an address range because some systems may be configured to ignore ECHO requests and not respond, and some firewalls and gateways may be configured to prevent ECHO requests from being transmitted to the destination network. Despite this shortcoming, Address sweeps are still used because they are simple to perform with software tools that automate the process.

Use the `icmp_sweep` anomaly in a DoS policy to protect against address sweeps.

There are a number of IPS signatures to detect the use of ICMP probes that can gather information about your network. These signatures include `AddressMask`, `Traceroute`, `ICMP.Invalid.Packet.Size`, and `ICMP.Oversized.Packet`. Include ICMP protocol signatures in your IPS sensors to protect against these probes/attacks.

Port scans

Potential attackers may run a port scan on one or more of your hosts. This involves trying to establish a communication session to each port on a host. If the connection is successful, a service may be available that the attacker can exploit.

Use the DoS anomaly check for `tcp_port_scan` to limit the number of sessions (complete and incomplete) from a single source IP address to the configured threshold. If the number of sessions exceed the threshold, the configured action is taken.

Use the DoS anomaly check for `udp_scan` to limit UDP sessions in the same way.

Probes using IP traffic options

Every TCP packet has space reserved for eight flags or control bits. They are used for communicating various control messages. Although space in the packet is reserved for all eight, there are various combinations of flags that should never happen in normal network operation. For example, the SYN flag, used to initiate a session, and the FIN flag, used to end a session, should never be set in the same packet.

Attackers may create packets with these invalid combinations to test how a host will react. Various operating systems and hardware react in different ways, giving a potential attackers clues about the components of your network.

The IPS signature `TCP.Bad.Flags` detects these invalid combinations. The default action is pass though you can override the default and set it to **Block** in your IPS sensor.

Configure packet replay and TCP sequence checking

The anti-replay CLI command allows you to set the level of checking for packet replay and TCP sequence checking (or TCP Sequence (SEQ) number checking). All TCP packets contain a Sequence Number (SEQ) and an Acknowledgement Number (ACK). The TCP protocol uses these numbers for error free end-to-end communications. TCP sequence checking can also be used to validate individual packets.

FortiGate units use TCP sequence checking to make sure that a packet is part of a TCP session. By default, if a packet is received with sequence numbers that fall out of the expected range, the FortiGate unit drops the packet. This is normally a desired behavior, since it means that the packet is invalid. But in some cases you may want to configure different levels of anti-replay checking if some of your network equipment uses non-RFC methods when sending packets.

Configure the anti-replay CLI command:

```
config system global
    set anti-replay {disable | loose | strict}
```

```
end
```

You can set anti-replay protection to the following settings:

- `disable` — No anti-replay protection.
- `loose` — Perform packet sequence checking and ICMP anti-replay checking with the following criteria:
 - The SYN, FIN, and RST bit can not appear in the same packet.
 - The FortiGate unit does not allow more than one ICMP error packet through before it receives a normal TCP or UDP packet.
 - If the FortiGate unit receives an RST packet, and `check-reset-range` is set to `strict`, the FortiGate unit checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- `strict` — Performs all of the loose checking but for each new session also checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value for each new session. Strict anti-replay checking can also help prevent SYN flooding.

If any packet fails a check it is dropped.

Configure ICMP error message verification

Enable ICMP error message verification to ensure an attacker can not send an invalid ICMP error message.

```
config system global
    check-reset-range {disable | strict}
end
```

- `disable` — the FortiGate unit does not validate ICMP error messages.
- `strict` — enable ICMP error message checking.

If the FortiGate unit receives an ICMP error packet that contains an embedded IP(A,B) | TCP(C,D) header, then if FortiOS can locate the A:C->B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped. Strict checking also affects how the anti-replay option checks packets.

Protocol header checking

Select the level of checking performed on protocol headers.

```
config system global
    check-protocol-header {loose | strict}
end
```

- `loose` — the FortiGate unit performs basic header checking to verify that a packet is part of a session and should be processed. Basic header checking includes verifying that the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options are correct, etc.
- `strict` — the FortiGate unit does the same checking as above plus it verifies that ESP packets have the correct sequence number, SPI, and data length.

If the packet fails header checking it is dropped by the FortiGate unit.

Evasion techniques

Attackers employ a wide range of tactics to try to disguise their techniques. If an attacker disguises a known attack in such a way that it is not recognized, the attack will evade your security and possibly succeed. FortiGate security recognizes a wide variety of evasion techniques and normalizes data traffic before inspecting it.

Packet fragmentation

Information sent across local networks and the Internet is encapsulated in packets. There is a maximum allowable size for packets and this maximum size varies depending on network configuration and equipment limitations. If a packet arrives at a switch or gateway and it is too large, the data it carries is divided among two or more smaller packets before being forwarded. This is called fragmentation.

When fragmented packets arrive at their destination, they are reassembled and read. If the fragments do not arrive together, they must be held until all of the fragments arrive. Reassembly of a packet requires all of the fragments.

The FortiGate unit automatically reassembles fragmented packets before processing them because fragmented packets can evade security measures. This reassembly of packets affects TCP, UDP and IP packets. There can be some variation though in what process does the reassembling. The IPS engine, nTurbo and the kernel all can do defragmentation.

For example, you have configured the FortiGate unit to block access to the example.org web site. Any checks for example.com will fail if a fragmented packet arrives and one fragment contains `http://www.exa` while the other contains `mple.com/`. Viruses and malware can be fragmented and avoid detection in the same way. The FortiGate unit will reassemble fragmented packets before examining network data to ensure that inadvertent or deliberate packet fragmentation does not hide threats in network traffic.

Non-standard ports

Most traffic is sent on a standard port based on the traffic type. The FortiGate unit recognizes most traffic by packet content rather than the TCP/UDP port and uses the proper IPS signatures to examine it. Protocols recognized regardless of port include DHCP, DNP3, FTP, HTTP, IMAP, MS RPC, NNTP, POP3, RSTP, SIP, SMTP, and SSL, as well as the supported IM/P2P application protocols.

In this way, the FortiGate unit will recognize HTTP traffic being sent on port 25 as HTTP rather than SMTP, for example. Because the protocol is correctly identified, the FortiGate unit will examine the traffic for any enabled HTTP signatures.

Negotiation codes

Telnet and FTP servers and clients support the use of negotiation information to allow the server to report what features it supports. This information has been used to exploit vulnerable servers. To avoid this problem, the FortiGate unit removes negotiation codes before IPS inspection.

HTTP URL obfuscation

Attackers encode HTML links using various formats to evade detection and bypass security measures. For example, the URL `www.example.com/cgi.bin` could be encoded in a number of ways to avoid detection but still work properly, and be interpreted the same, in a web browser.

The FortiGate prevents the obfuscation by converting the URL to ASCII before inspection.

HTTP URL obfuscation types

Encoding type	Example
No encoding	<code>http://www.example.com/cgi.bin/</code>

Encoding type	Example
Decimal encoding	http://www.example.com/c g i . b i n /
URL encoding	http://www.example.com/%43%47%49%2E%42%49%4E%2F
ANSI encoding	http://www.example.com/%u0063%u0067%u0069%u002E%u0062%u0069%u006E/
Directory traversal	http://www.example.com/cgi.bin/test/..

HTTP header obfuscation

The headers of HTTP requests or responses can be modified to make the discovery of patterns and attacks more difficult. To prevent this, the FortiGate unit will:

- remove junk header lines
- reassemble an HTTP header that's been folded onto multiple lines
- move request parameters to HTTP POST body from the URL

The message is scanned for any enabled HTTP IPS signatures once these problems are corrected.

HTTP body obfuscation

The body content of HTTP traffic can be hidden in an attempt to circumvent security scanning. HTTP content can be GZipped or deflated to prevent security inspection. The FortiGate unit will uncompress the traffic before inspecting it.

Another way to hide the contents of HTTP traffic is to send the HTTP body in small pieces, splitting signature matches across two separate pieces of the HTTP body. The FortiGate unit reassembles these 'chunked bodies' before inspection.

Microsoft RPC evasion

Because of its complexity, the Microsoft Remote Procedure Call protocol suite is subject to a number of known evasion techniques, including:

- SMB-level fragmentation
- DCERPC-level fragmentation
- DCERPC multi-part fragmentation
- DCERPC UDP fragmentation
- Multiple DCERPC fragments in one packet

The FortiGate unit reassembles the fragments into their original form before inspection.

Defending against DoS attacks

A denial of service is the result of an attacker sending an abnormally large amount of network traffic to a target system. Having to deal with the traffic flood slows down or disables the target system so that legitimate users can not use it for the duration of the attack.

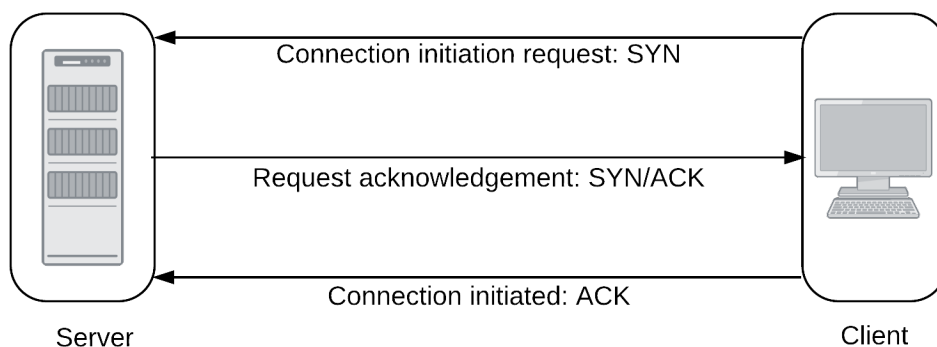
Any network traffic the target system receives has to be examined, and then accepted or rejected. TCP, UDP, and ICMP traffic is most commonly used, but a particular type of TCP traffic is the most effective. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are started between systems.

The “three-way handshake”

Communication sessions between systems start with establishing a TCP/IP connection. This is a simple three step process, sometimes called a “three-way handshake,” initiated by the client attempting to open the connection.

1. The client sends a TCP packet with the SYN flag set. With the SYN packet, the client informs the server of its intention to establish a connection.
2. If the server is able to accept the connection to the client, it sends a packet with the SYN and the ACK flags set. This simultaneously acknowledges the SYN packet the server has received, and informs the client that the server intends to establish a connection.
3. To acknowledge receipt of the packet and establish the connection, the client sends an ACK packet.

Establishing a TCP/IP connection



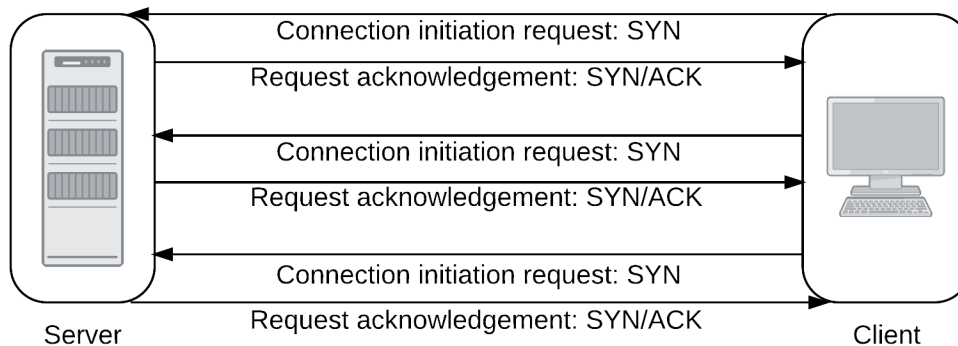
The three-way handshake is a simple way for the server and client to each agree to establish a connection and acknowledge the other party expressing its intent. Unfortunately, the three-way handshake can be used to interfere with communication rather than facilitate it.

SYN flood

When a client sends a SYN packet to a server, the server creates an entry in its session table to keep track of the connection. The server then sends a SYN+ACK packet expecting an ACK reply and the establishment of a connection.

An attacker intending to disrupt a server with a denial of service (DoS) attack can send a flood of SYN packets and not respond to the SYN+ACK packets the server sends in response. Networks can be slow and packets can get lost so the server will continue to send SYN+ACK packets until it gives up, and removes the failed session from the session table. If an attacker sends enough SYN packets to the server, the session table will fill completely, and further connection attempts will be denied until the incomplete sessions time out. Until this happens, the server is unavailable to service legitimate connection requests.

A single client launches a SYN flood attack




SYN floods are seldom launched from a single address so limiting the number of connection attempts from a single IP address is not usually effective.

SYN spoofing

With a flood of SYN packets coming from a single attacker, you can limit the number of connection attempts from the source IP address or block the attacker entirely. To prevent this simple defense from working, or to disguise the source of the attack, the attacker may spoof the source address and use a number of IP addresses to give the appearance of a distributed denial of service (DDoS) attack. When the server receives the spoofed SYN packets, the SYN+ACK replies will go to the spoofed source IP addresses which will either be invalid, or the system receiving the reply will not know what to do with it.

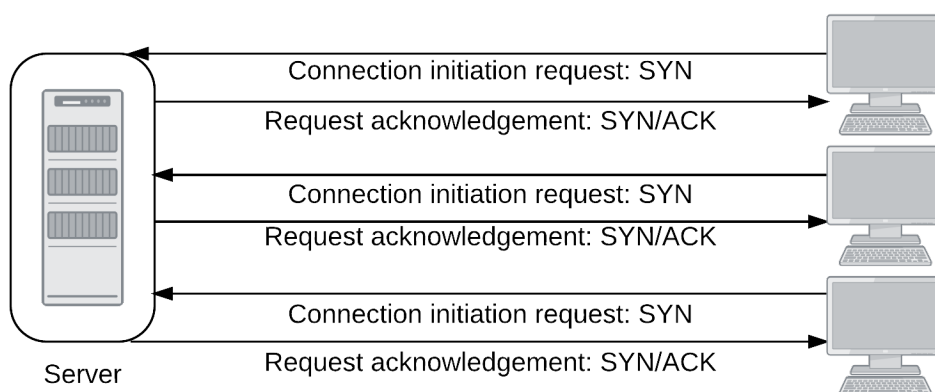
A client launches a SYN spoof attack



DDoS SYN flood

The most severe form of SYN attack is the distributed SYN flood, one variety of distributed denial of service attack (DDoS). Like the SYN flood, the target receives a flood of SYN packets and the ACK+SYN replies are never answered. The attack is distributed across multiple sources sending SYN packets in a coordinated attack.

Multiple attackers launch a distributed SYN flood



The distributed SYN flood is more difficult to defend against because multiple clients are capable of creating a larger volume of SYN packets than a single client. Even if the server can cope, the volume of traffic may

overwhelm a point in the network upstream of the targeted server. The only defence against this is more bandwidth to prevent any choke-points.

Configuring the SYN threshold to prevent SYN floods

The preferred primary defence against any type of SYN flood is the DoS anomaly check for `tcp_syn_flood` threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to **Pass**, the FortiGate unit will allow the SYN packets that exceed the threshold. If the action is set to **Block**, the FortiGate unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet when a SYN+ACK response is not received from the server. These tools will not “retry.” Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to **Block**.

SYN proxy

FortiGate units with network acceleration hardware, whether built-in or installed in the form of an add-on module, offer a third action for the `tcp_syn_flood` threshold. Instead of **Block** and **Pass**, you can choose to **Proxy** the incomplete connections that exceed the threshold value.

When the `tcp_syn_flood` threshold action is set to **f**, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the FortiGate unit receives an ACK response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

Other flood types

UDP and ICMP packets can also be used for DoS attacks, though they are less common. TCP SYN packets are so effective because the target receives them and maintains a session table entry for each until they time out. Attacks using UDP or ICMP packets do not require the same level of attention from a target, rendering them less effective. The target will usually drop the offending packets immediately, closing the session.

Use the `udp_flood` and `icmp_flood` thresholds to defend against these DoS attacks.

DoS policies

DDoS attacks vary in nature and intensity. Attacks aimed at saturating the available bandwidth upstream of your service can only be countered by adding more bandwidth. DoS policies can help protect against DDoS attacks that aim to overwhelm your server resources.

DoS policy recommendations

- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received.
- It is important to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, although your network may require very different values.


- One way to find the correct values for your environment is to set the action to **Pass** and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

Inside FortiOS: Denial of Service (DoS) Protection

FortiOS DoS protection maintains network integrity and performance by identifying and blocking harmful IPv4 and IPv6-based denial of service (DoS) attacks.

About DoS and DDoS attacks

A denial of service (DoS) occurs when an attacker overwhelms server resources by flooding a target system with anomalous data packets, rendering it unable to service genuine users. A distributed denial of service (DDoS) occurs when an attacker uses a master computer to control a network of compromised systems, otherwise known as a 'botnet', which collectively inundates the target system with excessive anomalous data packets.

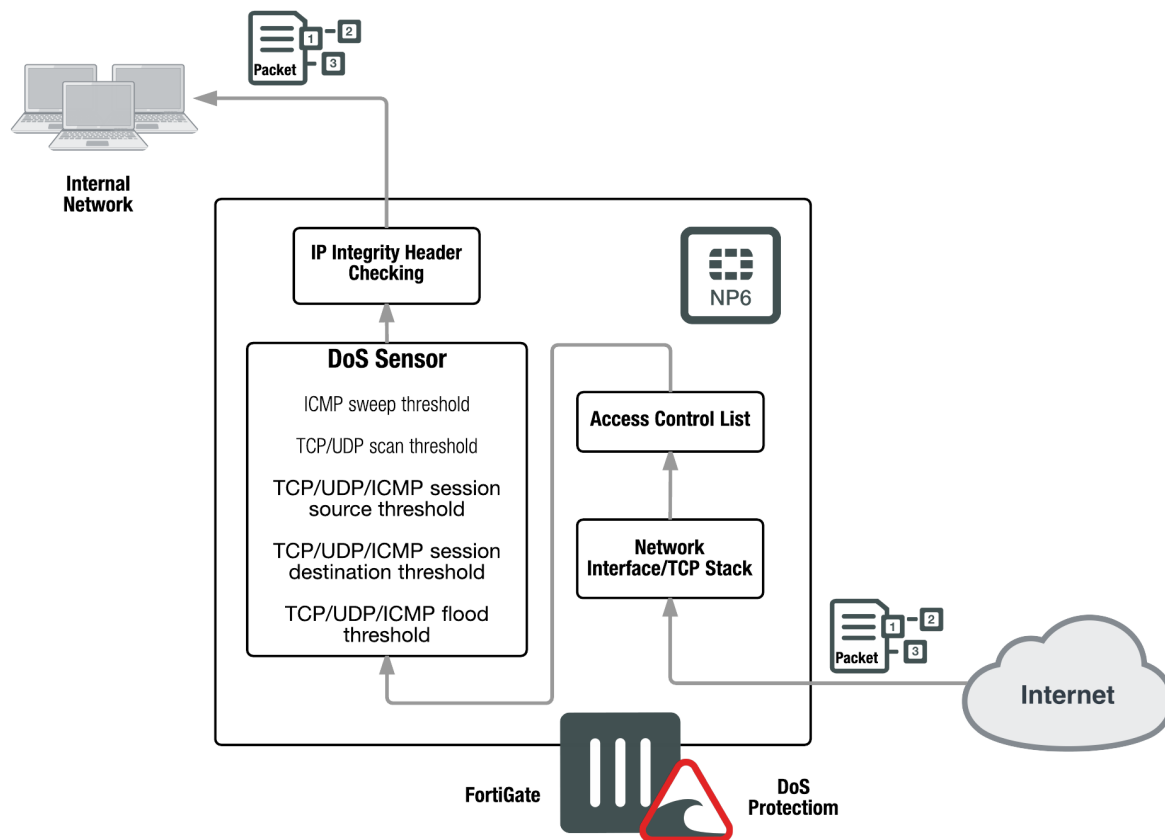


FortiOS DoS and DDoS protection

FortiOS DoS protection identifies potentially harmful traffic that could be part of a DoS or a DDoS attack by looking for specific traffic anomalies. Traffic anomalies that become DoS attacks include: TCP SYN floods, UDP floods, ICMP floods, TCP port scans, TCP session attacks, UDP session attacks, ICMP session attacks, and ICMP sweep attacks. Only traffic identified as part of a DoS attack is blocked; connections from legitimate users are processed normally.

FortiOS applies DoS protection very early in its traffic processing sequence to minimize the effect of a DoS attack on FortiOS system performance. DoS protection is the first step for packets after they are received by a FortiGate interface. Potential DoS attacks are detected and blocked before the packets are sent to other FortiOS systems.

FortiOS also includes an access control list feature that is implemented next. This accelerated ACL technology uses NP6 processors to block traffic (including DoS attacks) by source and destination address and service again before the packets are sent to the FortiGate CPU.



FortiOS DoS protection can operate in a standard configuration or operate out of band in sniffer mode, also known as one-arm mode, similar to intrusion detection systems. When operating in sniffer mode the FortiGate unit detects attacks and logs them without blocking them.

FortiOS DoS policies determine the course of action to take when anomalous traffic reaches a configured packet rate threshold. You can block an attacker, block an interface, block an attacker and interface, or allow traffic to pass through for monitoring purposes. This allows you to maintain network security by gathering information about attacks, monitor potentially offending traffic, or block offenders for the most protection.

FortiGates with NP6 processors also support synproxy DoS protection. An NP6-accelerated TCP SYN proxy offloads the three-way TCP handshake TCP SYN anomaly checking DoS protection to NP6 processors.

FortiOS DDoS Prevention

In addition to using DoS protection for protection against DoS attacks, FortiOS includes a number of features that prevent the spread of Botnet and C&C activity. Mobile Malware or Botnet and C&C protection keeps Botnet and C&C code from entering a protected network and compromising protected systems. As a result, systems on the protected network cannot become Botnet clients.


In addition, FortiOS can monitor and block outgoing Botnet connection attempts. Monitoring allows you to find and remove Botnet clients from your network and blocking prevents infected systems from communicating with Botnet sites.

Configuration options

Choose the standard configuration for maximum protection or configure sniffer mode to gather information.

Standard configuration

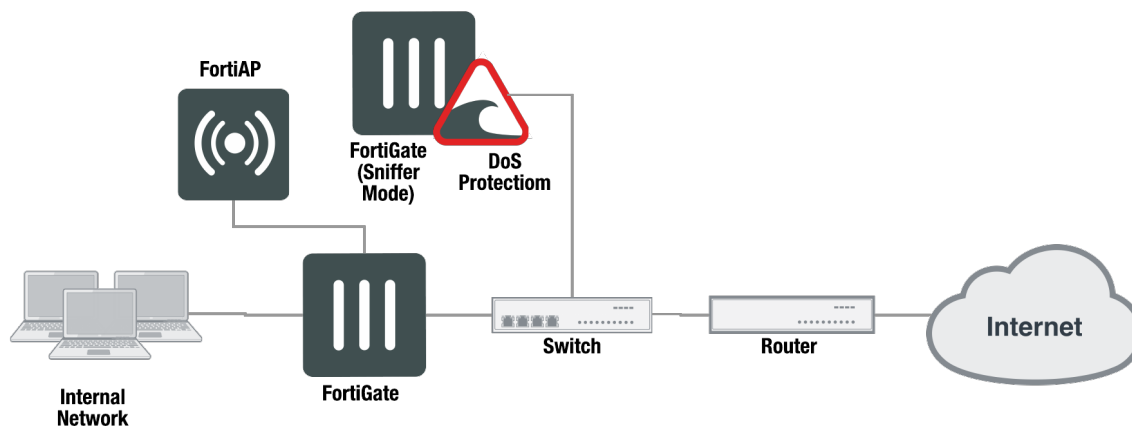
DoS protection is commonly configured on a FortiGate unit that connects a private or DMZ network to the Internet or on a FortiWiFi unit that connects a wireless LAN to an internal network and to the Internet. All Internet traffic or wireless LAN traffic passes through DoS protection in the FortiGate unit or the FortiWiFi unit.



Out of band configuration (sniffer mode)

A FortiGate unit in sniffer mode operates out of band as a one-armed Intrusion Detection System by detecting and reporting attacks. It does not process network traffic nor does it take action against threats. The FortiGate interface operating in sniffer mode is connected to a Test Access Point (TAP) or a Switch Port Analyzer (SPAN) port that processes all of the traffic to be analyzed. The TAP or SPAN sends a copy of the switch traffic to the out of band FortiGate for analysis.

FortiOS records log messages and sends alerts to system administrators when a DoS attack is detected. IDS scanning does not affect network performance or network traffic if the IDS fails or goes offline.



DoS policies

DoS policies provide effective and early DoS detection while remaining light on system resources. They are configured to monitor and to stop traffic with abnormal patterns or attributes. The DoS policy recognizes traffic as a threat when the traffic reaches a user-configured packet rate threshold. The policy then determines the appropriate action. In addition to choosing whether or not to log each type of anomaly, you can choose to pass or block threats.

DoS policy anomaly protection is applied to all incoming traffic to a single FortiGate interface, but you can narrow policies by specifying service, source address, and destination address. The FortiGate unit processes DoS policies in their own respective order first, followed by all other firewall policies.

Hardware acceleration

Hardware acceleration enhances protection and increases the efficiency of your network. FortiOS integrated Content Processors (CPs), Network Processors (NPs), and Security Processors (SPs) accelerate specialized security processing. DoS SYN proxy protection is built in to NP6 processors and many Fortinet Security Processors, like the CE4, XE2, and FE8, to guard against TCP SYN floods. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are initiated between systems. NP6 and SP processors can offload TCP SYN flood attack detection and blocking. The SP module increases a FortiGate unit's capacity to protect against TCP SYN flood attacks while minimizing the effect of attacks on the FortiGate unit's overall performance and the network performance. The result is improved capacity and overall system performance.

The FortiGuard Center

The FortiGuard Center shows information on all the most recent FortiGuard news, including information concerning zero-day research and hot intrusion detections. Research papers are also available that concern a

variety of current security issues.

To view recent developments, go to <http://www.fortiguard.com/static/intrusionprevention.html>.

Firewall Policies

The firewall policies of the FortiGate are one of the most important aspects of the appliance. There are a lot of building blocks and configurations involved in setting up a firewall and it within the policies that a lot of these components come together to form a cohesive unit to perform the firewall's main function, analyzing network traffic and responding appropriately to the results of that analysis.

There are a few different kinds of policies and in most cases these are further divided into IPv4 and IPv6 versions:

- [IPv4 Policy](#) - used for managing traffic going through the appliance using IPv4 protocols
- [IPv6 Policy](#) - used for managing traffic going through the appliance using IPv6 protocols
- [NAT64 Policy](#) - used for managing traffic going through the appliance that converts from IPv6 on the incoming interface to IPv4 on the outgoing interface
- [NAT46 Policy](#) - used for managing traffic going through the appliance that converts from IPv4 on the incoming interface to IPv6 on the outgoing interface
- [Multicast Policy](#) - used to manage traffic sent to multiple destinations
- [IPv4 Access Control List](#) - used to filter out packets based on specific IPV4 parameters.
- [IPv6 Access Control List](#) - used to filter out packets based on specific IPV6 parameters.
- [IPv4 DoS Policy](#) - used to prevent malicious or flawed packets on an IPv4 interface from denying access to users.
- [IPv6 DoS Policy](#) - used to prevent malicious or flawed packets on an IPv6 interface from denying access to users.

Because the policy determines whether or not NAT will be used, it is also important to look at how to configure:

- [Central SNAT](#) - used for granular controlling when NATing is in use.

IPv4 Policy

To configure a IPv4 policy in the GUI

1. Goto **Policy & Objects > IPv4 Policy**

The right side window will display a table of the existing IPv4 Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Make sure the policy has a name in the **Name** field



By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI, or in the GUI if you have first enabled the GUI option in the CLI.

- #### 3. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#)

4. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
5. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
8. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
9. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session
 - **LEARN** - collects information about the traffic for future analysis
 - **IPsec** - for using with IPsec tunnels

Because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Set the NAT parameter by toggling the slider button. (gray means it is disabled)

The NAT setting section is affected by whether or not Central NAT is enabled.

If Central NAT is enabled, the only option in Firewall / Network options will be whether to enable or disable NAT. The rest of the NAT parameters will be set in the Central SNAT page.

If Central NAT is disabled, there are two additional settings in the Policy configuration page.

11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the + icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "+" icon next to the **Search** field is a shortcut for creating a new IP Pool.

Security Profiles

13. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The "+" icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **DNS Filter**
- **Application Control**
- **CASI**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**
- **VoIP**
- **ICAP**
- **Web Application Firewall**
- **Proxy Options**
- **SSL/SSH Inspection**

Logging Options

14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the **Log Violation Traffic** setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

Settings if the LEARN action is selected

To get more information on the **LEARN** option, read the Learning mode for Firewall policies topic in [What's new for the Firewall in 5.4](#)

Firewall / Network Options

10. Set the **NAT** parameter by toggling the slider button. (gray means it is disabled). Unlike the **ACCEPT** option, whether or not Central NAT is enabled or disabled does not affect this settings options.
11. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
12. Toggle whether or not to **Enable this policy**. The default is enabled.
13. Select the **OK** button to save the policy.

Settings if the IPsec action is selected

VPN Tunnel

10. For the VPN Tunnel field, use the drop down menu to select the VPN tunnel that you want the policy associated with.
11. Toggle the sliding button to enable or disable the option to **Allow traffic to be initiated from the remote site**

Security Profiles

12. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The "+" icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **DNS Filter**
- **Application Control**
- **CASI**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**
- **VoIP**
- **ICAP**
- **Web Application Firewall**
- **Proxy Options**
- **SSL/SSH Inspection**

Logging Options

13. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
14. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
15. Toggle whether or not to **Enable this policy**. The default is enabled.
16. Select the **OK** button to save the policy.

IPv6 Policy

To configure a IPv6 policy in the GUI

1. Goto **Policy & Objects > IPv6 Policy**

The right side window will display a table of the existing IPv6 Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Make sure the policy has a name in the **Name** field



By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI, or in the GUI if you have first enabled the GUI option in the CLI.

3. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#)
4. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
5. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
8. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
9. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Set the NAT parameter by toggling the slider button. (gray means it is disabled)
11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:

- **Use Outgoing Interface Address**
- **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the **+** icon next to the **Search** field is a shortcut for creating a new IP Pool.

Security Profiles

13. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The **+** icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **Application Control**
- **CASI**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**
- **VoIP**
- **ICAP**
- **SSL/SSH Inspection**

Logging Options

14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the Log Violation Traffic setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

NAT64 Policy

To configure a NAT64 policy in the GUI

1. Goto **Policy & Objects > NAT64 Policy**

The right side window will display a table of the existing NAT64 Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#)
 3. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
 4. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
 7. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
 8. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Skip the NAT setting. This type of policy is intended only for traffic that is being NATed from IPv6 to IPv4, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the + icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "+" icon next to the **Search** field is a shortcut for creating a new IP Pool.

14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).

If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.

15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the **Log Violation Traffic** setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

NAT46 Policy

To configure a NAT46 policy in the GUI

1. Goto **Policy & Objects > NAT46 Policy**

The right side window will display a table of the existing NAT46 Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#)
 3. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
 4. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)

7. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
8. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Skip the NAT setting. This type of policy is intended only for traffic that is being NATed from IPv4 to IPv6, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the + icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "+" icon next to the **Search** field is a shortcut for creating a new IP Pool.
14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the **DENY** action is selected

Enable the Log Violation Traffic setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

Central SNAT

Central NAT is disabled by default. To toggle the feature on or off, use the following commands:

```
config system settings
  set central-nat [enable | disable]
end
```

When Central NAT is enable the **Central SNAT** section will appear under the Policy & Objects heading in the GUI.

The Central SNAT window contains a table of all of the Central SNAT policies.

To configure a Central SNAT entry in the GUI

1. Goto **Policy & Objects > Central SNAT**

The right side window will display a table of the existing Central SNAT entries.

- To edit an existing entry, double click on the policy you wish to edit
 - To create a new entry, select the **Create New** icon in the top left side of the right window.
2. Set the **Source Address** parameter by selecting an address from the drop down menu. One or more addresses can be selected. Additional addresses can be added later by selecting the circle icon with the "+" symbol inside it. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 3. Set the **Destination Address** parameter by selecting an address from the drop down menu. One or more addresses can be selected. Additional addresses can be added later by selecting the circle icon with the "+" symbol inside it.
 4. Set the **Translated Address** parameter by selecting an IP Pool from the drop down menu. There are four types of IP Pools. The type selected will determine which further settings are to be set.
 5. Set the **Protocol** parameter.

There are 5 options for the **Protocol**.

- **ANY** - any protocol traffic
 - **TCP** - TCP traffic only. Protocol number set to 6
 - **UDP** - UDP traffic only . Protocol number set to 17
 - **SCTP** - SCTP traffic only. Protocol number set to 132
 - **Specify** - User can specify the traffic filter protocol by setting the protocol number in the field.
6. If the IP Pool is of the type: Overload, **Explicit Port Mapping** can be enabled.
To enable or disable, use the check box. Once enabled, the following additional parameters will appear.
 - **Original Source Port** - in the left number field, set the starting number of the source port range.
 - **Translated Port** - in the left number field, set the starting number of the translated port range. If it is a single port range leave the right number field alone. If the right number field is set to a number higher than the left, the right number field for the Original Source Port will change to make sure the 2 number ranges have a matching number of ports.
 7. Select the **OK** button to save the entry.

To configure Central SNAT in the CLI

1. Using the CLI interface of your choice, run the following command to get to the correct context.

```
config firewall central-snat-map
```

 - To edit an existing entry, run the command `show` or `show full-configuration` to get a listing of all of the entries in the map. Take note of the policy ID fo the entry to be edited.
 - To create a new entry the next step will use the policy ID 0 which will check for an unused ID number and create an entry with that number.
2. Edit or create an entry with the correct policy ID

```
edit <policyID number>
```


Run the following commands to set the parameters of the entry:

```
set status enable
set orig-addr <valid address object preconfigured on the FortiGate>
set dst-addr <valid address object preconfigured on the FortiGate>
set nat-ippool <valid ippool object preconfigured on the FortiGate>
set protocol <integer for protocol number>
```

3. Save the entry by running the command `end` or `next`.

Example: Central NAT Table

The company has a server on the Development LAN that needs to communicate with a server at a remote site over the Internet. One of the restrictions on the communications between these systems is that the IP address and source port must be specific.

- The traffic going out on to the Internet must be NATed
- The traffic is coming from a server with the IP address 192.168.150.86
- An address called "app-server" has been created for the address 192.168.150.86 on the port1 interface
- An IP Pool named "Connection to Example.com" has been created that matches the internal IP address of 192.168.150.86 to the external address of 256.23.45.67
- An address called "app-server-ext" has been created for the address 256.23.45.67 on the wan1 interface
- The remote servers is set to accept traffic from 256.23.45.67
- The originating traffic from the server originates in the port range from 2000 to 3000
- The remote site requires that the source TCP port must be within the 12000 to 13000 range

The original address and Translated Address fields require values that are address names that are listed in the address section of Firewall Objects.

Configuring the address in the GUI

1. Go to **Policy & Objects > Policy > Central SNAT > Create New**.
2. Fill out the fields with the following information:

Source Address	app-server
Destination Address	app-server-ext
Translated Address	connection-example.com
Protocol	ANY
Explicit Port Mapping	<enabled>
Original Source Port	2000
Translated Port	12000-13000

3. Select **OK**.

Configuring the address in the CLI

Enter the following CLI commands:

```
config firewall central-snat-map
```

```
edit 0
  set status enable
  set orig-addr app-server
  set dst-addr app-svr-ext
  set nat-ippool connection-example.com
  set protocol 0
  set orig-port 2000
  set nat-port 12000-13000
end
```

To verify that the table was added correctly:

1. Go to **Policy & Objects > Policy > Central NAT**.
 2. Check that the table has been added to the list of Central NAT Tables and that the listed settings are correct.
- or

1. Enter the following CLI command:

```
config firewall central-nat
show full-configuration
```

2. Verify that the listing of tables includes the one that you have just configured, with the correct settings.



When configuring the Central NAT in the GUI you may notice that only those addresses which have been configured to be associated with **any** interface are displayed in the drop down menu for choosing a Source Address and yet the CLI will allow any address to be used, not just those associated with **any** interface. This is because by default the policies in the GUI use a function of cross referencing which addresses are allowed based on which interface is involved in the policy. When combined with the aspect of Central NAT that doesn't restrict to a specific interface. This means the only addresses will be allowed are those associated with the **any** interface. The CLI does not have this cross referencing function which is why the CLI seems less restrictive. However, more care must be taken when using the CLI to make sure that appropriate addresses are used.

IPv4 Access Control List

The **IPv4 Access Control List** is a specialized policy for denying IPv4 traffic based on:

- the incoming interface
- the source addresses of the traffic
- the destination addresses of the traffic
- the services or ports the traffic is using

The only action available in this policy is **DENY**

For more information on see [Access Control Lists](#)

To configure a IPv4 Access Control List entry in the GUI

1. Goto **Policy & Objects > IPv4 Access Control List**

The right side window will display a table of the existing IPv4 Access Control List entries.

- To edit an existing entry, double click on the policy you wish to edit
 - To create a new entry, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
 3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
 6. Toggle whether or not to **Enable this policy**. The default is enabled.
 7. Select the **OK** button to save the policy.

To configure a IPv4 Access Control List entry in the CLI

Use the following syntax:

```
config firewall acl
  edit <acl Policy ID #>
    set status enable
    set interface <interface>
    set srcaddr <address object>
    set dstaddr <address object>
    set service <service object>
  end
end
```

IPv6 Access Control List

The **IPv6 Access Control List** is a specialized policy for denying IPv6 traffic based on:

- the incoming interface
- the source addresses of the traffic
- the destination addresses of the traffic
- the services or ports the traffic is using

The only action available in this policy is **DENY**

To configure a IPv6 Access Control List entry in the GUI

1. Goto **Policy & Objects > IPv6 Access Control List**

The right side window will display a table of the existing IPv4 Access Control List entries.

- To edit an existing entry, double click on the policy you wish to edit
- To create a new entry, select the **Create New** icon in the top left side of the right window.

2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
6. Toggle whether or not to **Enable this policy**. The default is enabled.
7. Select the **OK** button to save the policy.

To configure a IPv6 Access Control List entry in the CLI

Use the following syntax:

```
config firewall acl6
  edit <acl Policy ID #>
    set status enable
    set interface <interface>
    set srcaddr <address object>
    set dstaddr <address object>
    set service <service object>
  end
end
```

IPv4 DoS Policy

To configure a IPv4 DoS Policy in the GUI

1. Goto **Policy & Objects > IPv4 DoS Policy**

The right side window will display a table of the existing IPv4 DoS Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
 3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
 6. Set the parameters for the various traffic anomalies.

All of the anomalies that profiles have been created for are in 2 tables. These tables break up the anomaly profiles into **L3 Anomalies** and **L4 Anomalies**. All of the anomalies have the following parameters that can be set on a per anomaly or per column basis.

- Status - enable or disable the indicated profile
- Logging - enable or disable logging of the indicated profile being triggered
- Action - whether to Pass or Block traffic when the threshold is reached
- Threshold - the number of anomalous packets detected before triggering the action.

The listing of anomaly profiles includes:

L3 Anomalies

- ip_src_session
- ip_dst_session

L4 Anomalies

- tcp_syn_flood
- tcp_port_scan
- tcp_src_session
- tcp_dst_session
- udp_flood
- udp_scan
- udp_src_session
- udp_dst_session
- icmp_flood
- icmp_sweep
- icmp_src_session
- sctp_flood
- sctp_scan
- sctp_src_session
- sctp_dst_session

7. Toggle whether or not to **Enable this policy**. The default is enabled.
8. Select the **OK** button to save the policy.

Example

The company wishes to protect against Denial of Service attack. They have chosen some where they wish to block the attacks of the incidence goes above a certain threshold and for some others they are just trying to get a baseline of activity for those types of attacks so they are letting the traffic pass through without action.

- The interface to the Internet is on WAN1
- There is no requirement to specify which addresses are being protected or protected from.
- The protection is to extend to all services.
- The TCP attacks are to be blocked

- The UDP, ICMP, and IP attacks are to be recorded but not blocked.
- The SCTP attack filters are disabled
- The tcp_syn_flood attach's threshold is to be changed from the default to 1000

Configuring the DoS Policy in the GUI

1. Go to **Policy & Objects > Policy > DoS**.
2. Create a new policy
3. Fill out the fields with the following information:

Field	Value
Incoming Interface	wan1
Source Address	all
Destination Addresses	all
Service	ALL

L3 Anomalies

Name	Status	Logging	Action	Threshold
ip_src_session	enabled	enabled	Pass	5000
ip_dst_session	enabled	enabled	Pass	5000

L4 Anomalies

Name	Status	Logging	Action	Threshold
tcp_syn_flood	enabled	enabled	Block	1000
tcp_port_scan	enabled	enabled	Block	<default value>
tcp_src_session	enabled	enabled	Block	<default value>
tcp_dst_session	enabled	enabled	Block	<default value>
udp_flood	enabled	enabled	Pass	<default value>
udp_scan	enabled	enabled	Pass	<default value>
udp_src_session	enabled	enabled	Pass	<default value>
udp_dst_session	enabled	enabled	Pass	<default value>

Name	Status	Logging	Action	Threshold
icmp_flood	enabled	enabled	Pass	<default value>
icmp_sweep	enabled	enabled	Pass	<default value>
icmp_src_session	enabled	enabled	Pass	<default value>
icmp_dst_session	enabled	enabled	Pass	<default value>
sctp_flood	not enabled	not enabled	Pass	<default value>
sctp_scan	not enabled	not enabled	Pass	<default value>
sctp_src_session	not enabled	not enabled	Pass	<default value>
sctp_dst_session	not enabled	not enabled	Pass	<default value>

4. Toggle the button next to **Enable this policy** to **ON**.
5. Select **OK**.

Configuring the DoS Policy in the GUI

Using the CLI of your choice, enter the following commands:

```
config firewall DoS-policy
edit 0
    set status enable
    set interface wan1
    set srcaddr all
    set dstaddr all
    set service ALL
    config anomaly
        edit "tcp_syn_flood"
            set status enable
            set log disable
            set action block
            set threshold 1000
        next
        edit "tcp_port_scan"
            set status enable
            set log disable
            set action block
            set threshold 1000
        next
        edit "tcp_src_session"
            set status enable
            set log disable
            set action block
            set threshold 5000
        next
        edit "tcp_dst_session"
            set status enable
            set log disable
            set action block
```

```
        set threshold 5000
    next
edit "udp_flood"
    set status enable
    set log disable
    set action pass
    set threshold 2000
    next
edit "udp_scan"
    set status enable
    set log disable
    set action pass
    set quarantine none
    set threshold 2000
    next
edit "udp_src_session"
    set status enable
    set log disable
    set action pass
    set threshold 5000
    next
edit "udp_dst_session"
    set status enable
    set log disable
    set action pass
    set threshold 5000
    next
edit "icmp_flood"
    set status enable
    set log disable
    set action pass
    set threshold 250
    next
edit "icmp_sweep"
    set status enable
    set log disable
    set action pass
    set threshold 100
    next
edit "icmp_src_session"
    set status enable
    set log disable
    set action pass
    set threshold 300
    next
edit "icmp_dst_session"
    set status enable
    set log disable
    set action pass
    set threshold 1000
    next
edit "ip_src_session"
    set status disable
    set log enable
    set action pass
    set threshold 5000
    next
```



```
edit "ip_dst_session"
    set status disable
    set log enable
    set action pass
    set threshold 5000
next
edit "sctp_flood"
    set status disable
    set log disable
    set action pass
    set threshold 2000
next
edit "sctp_scan"
    set status disable
    set log disable
    set action pass
    set threshold 1000
next
edit "sctp_src_session"
    set status disable
    set log disable
    set action pass
    set threshold 5000
next
edit "sctp_dst_session"
    set status disable
    set log disable
    set action pass
    set threshold 5000
next
end
end
end
```



In this example of the CLI, all of the relevant settings have been left in, but some of them are default settings and would not have to have been specifically set to work. For instance, if the action parameter is not set it automatically defaults to pass.

IPv6 DoS Policy

To configure a IPv6 DoS Policy in the GUI

1. Goto **Policy & Objects > IPv6 DoS Policy**

The right side window will display a table of the existing IPv6 DoS Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
3. Set the **Source IPv6 Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
4. Set the **Destination IPv6Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#)
6. Set the parameters for the various traffic anomalies.

All of the anomalies that profiles have been created for are in 2 tables. These tables break up the anomaly profiles into **L3 Anomalies** and **L4 Anomalies**. All of the anomalies have the following parameters that can be set on a per anomaly or per column basis.

- Status - enable or disable the indicated profile
- Logging - enable or disable logging of the indicated profile being triggered
- Action - whether to Pass or Block traffic when the threshold is reached
- Threshold - the number of anomalous packets detected before triggering the action.

The listing of anomaly profiles includes:

L3 Anomalies

- ip_src_session
- ip_dst_session

L4 Anomalies

- tcp_syn_flood
- tcp_port_scan
- tcp_src_session
- tcp_dst_session
- udp_flood
- udp_scan
- udp_src_session
- udp_dst_session
- icmp_flood
- icmp_sweep
- icmp_src_session
- icmp_dst_session
- sctp_flood
- sctp_scan

7. Toggle whether or not to **Enable this policy**. The default is enabled.
8. Select the **OK** button to save the policy.

Firewall objects

As was mentioned earlier, the components of the FortiGate firewall go together like interlocking building blocks. The Firewall objects are a prime example of those building blocks. They are something that can be configured once and then used over and over again to build what you need. They can assist in making the administration of the FortiGate unit easier and more intuitive as well as easier to change. By configuring these objects with their future use in mind as well as building in accurate descriptions the firewall will become almost self documenting. That way, months later when a situation changes, you can take a look at a policy that needs to change and use a different firewall object to adapt to the new situation rather than build everything new from the ground up to accommodate the change.

This chapter includes information about the following Firewall objects:

- Addresses
- Services and TCP ports
- Firewall schedules
- Security profiles

UUID Support

A Universally Unique Identified (UUID) attribute has been added to some firewall objects, so that the logs can record these UUID to be used by a FortiManager or FortiAnalyzer unit. The objects currently include:

- Addresses, both IPv4 and IPv6
- Address Groups, both IPv4 and IPv6
- Virtual IPs, both IPv4 and IPv6
- Virtual IP groups, both IPv4 and IPv6
- Policies, IPv4, IPv6 and IP64

A UUID is a 16-octet (128-bit) number that is represented by 32 lowercase hexadecimal digits. The digits are displayed in five groups separated by hyphens (-). The pattern is 8-4-4-4-12; 36 digits if you include the hyphens.



Note: UUID is only supported on large-partition platforms ($\geq 128M$)

Addresses

Firewall addresses define sources and destinations of network traffic and are used when creating policies. When properly set up these firewall objects can be used with great flexibility to make the configuration of firewall policies simpler and more intuitive. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

The address categories and the types within those categories on the FortiGate unit can include:

- IPv4 addresses
 - IP address and Netmask
 - IP address range
 - Geography based address
 - Fully Qualified Domain Name (FQDN) address
 - Wildcard FQDN
 - IPv4 Address Group
- IPv6 addresses
 - Subnets
 - IP range
 - IPv6 Address Group
- Multicast addresses
 - Multicast IP range
 - Broadcast subnets
- Explicit Proxy Address
 - URL Pattern
 - Host Regex Match
 - URL Category
 - HttpMethod
 - User Agent
 - HTTP Header
 - Advanced (Source)
 - Advanced (Destination)
- IP Pools (IPv4)
 - Overload
 - One-to-one
 - Fixed Port Range
 - Port Block Allocation
- IP Pools (IPv6)
- Virtual IP Addresses
 - IPv4
 - IPv6

- NAT46
- NAT64

Interfaces

When setting up an address one of the parameters that is asked for is the interface. This means that the system will expect to see that address only on the interface that you select. You can only select one interface. If you expect that the address may be seen at more than one interface you can choose the “any” interface option. Whenever, possible it is best to choose a more specific interface than the “any” option because in the GUI configuration of firewall policies there is a drop down field that will show the possible addresses that can be used. The drop down will only show those addresses that can be on the interface assigned for that interface in the policy.

Example:

- You have an address called “XYZ”.
- “XYZ” is set to the WAN1 interface because that is the only interface that will be able to access that address.
- When you are selecting a Source Address in the Web-based Manager for a policy that is using the DMZ the address “XYZ” will not be in the drop-down menu.

When there are only 10 or 20 addresses this is not a concern, but if there are a few hundred addresses configured it can make your life easier.

Addresses, address groups, and virtual IPs must have unique names to avoid confusion in firewall policies. If an address is selected in a policy, the address cannot be deleted until it is deselected from the policy.

Addressing Best Practices Tip



The other reason to assign a specific interface to addresses is that it will prevent you from accidentally assigning an address where it will not work properly. Using the example from earlier, if the “XYZ” address was assigned to the “Any” interface instead of WAN1 and you configure the “XYZ” address.

Addressing Best Practices Tip



Don't specify an interface for VIP objects or other address objects that may need to be moved or approached from a different direction. When configuring a VIP you may think that it will only be associated with a single interface, but you may later find that you need to reference it on another interface.

Example: Some web applications require the use of a FQDN rather than an IP address. If you have a VIP set up that works from the Internet to the Internal LAN you won't be able to use that VIP object to access it from an internal LAN interface.

IPv4 Addresses

When creating an IPv4 address there are a number of different types of addresses that can be specified. These include:

- FQDN
- Geography
- IP Range
- IP/Netmask
- Wildcard FQDN

Which one chosen will depend on which method most easily yet accurately describes the addresses that you are trying to include with as few entries as possible based on the information that you have. For instance, if you are trying to describe the addresses of a specific company's web server but if you have no idea of how extensive there web server farm is you would be more likely to use a Fully Qualified Domain Name (FQDN) rather than a specific IP address. On the other hand some computers don't have FQDNs and a specific IP address must be used.

The following is a more comprehensive description of the different types of addresses.

FQDN Addresses

By using Fully Qualified Domain Name (FQDN) addressing you can take advantage of the dynamic ability of DNS to keep up with address changes without having to manually change the addresses on the FortiGate. FQDN addresses are most often used with external web sites but they can be used for internal web sites as well if there is a trusted DNS server that can be accessed. FQDN addressing also comes in handy for large web sites that may use multiple addresses and load balancers for their web sites. The FortiGate firewall automatically maintains a cached record of all the addresses resolved by the DNS for the FQDN addresses used.

For example, if you were doing this manually and you wanted to have a security policy that involved Google you could track down all of the IP addresses that they use across multiple countries. Using the FQDN address is simpler and more convenient.

When representing hosts by an FQDN, the domain name can also be a subdomain, such as mail.example.com.

Valid FQDN formats include:

- <host_name>.<top_level_domain_name> such as example.com
- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as mail.example.com

When creating FQDN entries it is important to remember that:

- Wildcards are not supported in FQDN address objects
- While there is a level of convention that would imply it, "www.example.com" is not necessarily the same address of "example.com". they will each have their own records on the DNS server.

The FortiGate firewall keeps track of the DNS TTLs so as the entries change on the DNS servers the IP address will effectively be updated for the FortiGate. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache.



There is a possible security downside to using FQDN addresses. Using a fully qualified domain name in a security policy means that your policies are relying on the DNS server to be accurate and correct. DNS servers in the past were not seen as potential targets because the thinking was that there was little of value on them and therefore are often not as well protected as some other network resources. People are becoming more aware that the value of the DNS server is that in many ways it controls where users and computers go on the Internet. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **FQDN** from the drop down menu.
6. Input the domain name in the **FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: FQDN address

You have to great a policy that will govern traffic that goes to a site that has a number of servers on the Internet. Depending on the traffic or the possibility that one of the servers is down network traffic can go to any one of those sites. The consistent factor is that they all use the same Fully Qualified Domain Name.

- The FQDN of the web site: example.com
- The number of ISP connections off of the FortiGate firewall: 2

Configuring the address in the GUI

1. Go to **Policy & Objects> Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information:

Category	Address
Name	BigWebsite.com
Type	FQDN
FQDN	bigwebsite.com
Interface	any
Show in Address List	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

Configuring the address in the CLI

```
config firewall address
edit BigWebsite.com
set type fqdn
set associated-interface any
set fqdn bigwebsite.com
end
```

Verification

To verify that the addresses were added correctly:

1. Go to **Firewall Objects > Address > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Changing the TTL of a FQDN address

To make sure that the FQDN resolves to the most recent active server you have been asked to make sure that the FortiGate has not cached the address for any longer than 10 minutes.

There is no field for the cached time-to-live in the web-based manager. It is only configurable in the CLI. Enter the following commands:

```
config firewall address
edit BigWebsite.com
set cache-ttl 600
end
```

Geography Based Addresses

Geography addresses are those determined by country of origin.

This type of address is only available in the IPv4 address category.

Creating a Geography address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **Geography** from the drop down menu.
6. In the **Country** field, select a single country from the drop down menu.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.

8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: Geography-based Address

Configuring the address in the GUI

Your company is US based and has information on its web site that may be considered information that is not allowed to be sent to embargoed countries. In an effort to help reduce the possibility of sensitive information going to those countries you have been asked to set up addresses for those countries so that they can be blocked in the firewall policies.

- One of the countries you have been asked to block is Cuba
- You have been asked to comment the addresses so that other administrators will know why they have been created

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	Address
Name	Cuba
Type	Geography
Country	Cuba
Interface	any
Visibility	<enable>
Comments	Embargoed

3. Select **OK**.

Configuring the address in the CLI

Enter the following CLI commands:

```
config firewall address
edit Cuba
set type geography
set country CN
set interface wan1
end
```

Overrides

It is possible to assign a specific IP address range to a customized country ID. Generally, geographic addressing is done at the VDOM level; it could be considered global if you are using the root VDOM, but the geoip-override setting is a global setting.

```
config system geoip-override
```

```
edit "test"
  set country-id "A0"
  config ip-range
    edit 1
      set start-ip 7.7.7.7
      set end-ip 7.7.7.8
    next
  edit 2
    set start-ip 7.7.10.1
    set end-ip 7.7.10.255
  end
```



- While the setting exists in the configuration file, the system assigns the country-id option automatically.
- While you can use "edit 1" and "edit 2", it is simpler to use "edit 0" and let the system automatically assign an ID number.

After creating a customized Country by using geoip-override command, the New country name has been added automatically to the country list and will be available on the Firewall Address Country field.

Diagnose commands

There are a few diagnose commands used with geographic addresses. The basic syntax is:

```
diagnose firewall ipgeo [country-list | ip-list | ip2country | override |  
copyright-notice]
```

Diagnose command	Description
country-list	Listing of all the countries.
ip-list	List of the IP addresses associated with the country
ip2country	Used to determine which country a specific IP address is assigned to.
override	Listing of user defined geography data - items configured by using "config system geoip-override" command.
copyright-notice	Shows the copyright notice.



Click on the diagnose command in the table to connect to the Fortinet Diagnose Wiki page that deals with the command option, to get more information.

IP Range Addresses

Where the Subnet address is good at representing a standardized group of addresses that are subnets the IP Range type of address can describe a group of addresses while being specific and granular. It does this by specifying a continuous set of IP addresses between one specific IP address and another. While it is most common that this range is with a subnet it is not a requirement. For instance, 192.168.1.0/24 and 192.168.2.0/24 would be 2 separate subnets but if you wanted to describe the top half of one and the bottom half of the other you could describe the range of 192.168.1.128-192.168.2.127. It's also a lot easier than trying to calculate the correct subnet mask.

The format would be:

x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120

There is a notation that is commonly used and accepted by some devices that follows the format:

x.x.x.[x-x], such as 192.168.110.[100-120]

This format is not recognized in FortiOS 5.2 as a valid IP Range.

Creating a IP Range address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, choose **Address**(IPv4 addresses) or **IPv6 Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **IP Range** from the drop down menu.
6. In the **Subnet / IP Range** field, enter the range of addresses in the following format: x.x.x.x-x.x.x.x (no spaces)
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu. (This setting is not available for IPv6 addresses)
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	Address or IPv6 Address
Name	Guest_users
Type	IP Range

Field	Value
Subnet / IP Range	192.168.100.200-192.168.100.240
Interface	Port1
Show in Address List	[on]
Comments	Computers on the 1st floor used by guests for Internet access.



IP Range addresses can be configured for both IPv4 and IPv6 addresses. The only differences in creating an IPv6 IP Range address is that you would choose IPv6 Address for the Category and the syntax of the address in the Subnet/IP Range field would be in the format of 2001:0db8:0000:0002:0:0:0:20-2001:0db8:0000:0004:0:0:0:20

IP / Netmask Addresses

The subnet type of address is expressed using a host address and a subnet mask. From a strictly mathematical stand point this is the most flexible of the types because the address can refer to as little one individual address or as many as all of the available addresses.

It is usually used when referring to your own internal addresses because you know what they are and they are usually administered in groups that are nicely differentiated along the lines of the old A, B, and C classes of IPv4 addresses. They are also addresses that are not likely to change with the changing of Internet Service Providers (ISP).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- A single host such as a single computer with the address 192.45.46.45
- A range of hosts such as all of the hosts on the subnet 192.45.46.1 to 192.45.46.255
- All hosts, represented by 0.0.0.0 which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- Netmask for a class A subnet of 16,777,214 usable addresses: 255.0.0.0, or /8
- Netmask for a class B subnet of 65,534 usable addresses: 255.255.0.0, or /16
- Netmask for a class C subnet of 254 usable addresses: 255.255.255.0, or /24
- Netmask for subnetted class C of 126 usable addresses: 255.255.255.128, or /25
- Netmask for subnetted class C of 62 usable addresses: 255.255.255.128, or /26
- Netmask for subnetted class C of 30 usable addresses: 255.255.255.128, or /27
- Netmask for subnetted class C of 14 usable addresses: 255.255.255.128, or /28
- Netmask for subnetted class C of 6 usable addresses: 255.255.255.128, or /29
- Netmask for subnetted class C of 2 usable addresses: 255.255.255.128, or /30

- Netmask for a single computer: 255.255.255.255, or /32
- Netmask used with 0.0.0.0 to include all IP addresses: 0.0.0.0, or /0

So for a single host or subnet the valid format of IP address and netmask could be either:

x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

or

x.x.x.x/x, such as 192.168.1.0/24

Static Route Configuration

A setting that is found in the IP/Netmask address type that is not found in the other address types is the enabling or disabling of **Static Route Configuration**. Enabling this feature includes the address in the listing of named addresses when setting up a static route.

To use in the GUI

1. Enable the **Static Route Configuration** in the address.
2. Go to **Network > Static Routes** and create a new route.
3. For a **Destination** type, choose **Named Address**.
4. Using the drop down menu, enter the name of the address object in the field just underneath the **Destination** type options.
5. Fill out the other information relevant to the route
6. Select the **OK** button

To enable in the CLI:

```
config firewall address
edit <address_name>
set allow-routing enable
end
```

Creating a Subnet address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **IP/Netmask** from the drop down menu.
6. In the **Subnet/IP Range** field, enter the address and subnet mask according to the format x.x.x.x/x.x.x.x or the short hand format of x.x.x.x/x
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Select the desired on/off toggle setting for **Static Route Configuration**.
10. Input any additional information in the **Comments** field.
11. Press **OK**.

Example

Example of a Subnet address for a database server on the DMZ:

Field	Value
Category	Address
Name	DB_server_1
Type	IP/Netmask
Subnet/IP Range	United States
Interface	any
Show in Address List	[on]
Static Route Configuration	[off]
Comments	

Wildcard FQDN

There are a number of companies that use secondary and even tertiary domain names or FQDNs for their websites. Wildcard FQDN addresses are to ease the administrative overhead in cases where this occurs. Sometimes its as simple as sites that still use www. as a prefix for their domain name. If you don't know whether or not the www is being used it's simpler to use a wildcard and include all of the possibilities whether it be example.com, www.example.com or even ftp.example.com.



Wildcard FQDN addresses do not resolve to a specific set of IP addresses in the same way that a normal FQDN addresss does. They are intended for use in SSL exemptions and should not be used as source or destination addresses in policies.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** fUncategorizedield, select **Wildcard FQDN**from the drop down menu.
6. Input the domain name in the **Wildcard FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a FQDN address for a remote FTP server used by Accounting team:

Field	Value
Category	Address
Name	Example.com_servers
Type	Wildcard FQDN
Wildcard FQDN	*.example.com
Interface	any
Show in Address List	[on]
Comments	Secondary and tertiary domain names for example.com

IPv6 Addresses

When creating an IPv6 address there are a number of different types of addresses that can be specified. These include:

- Subnet
- IP Range - the details of this type of address are the same as the IPv4 version of this type

The IPv6 addresses don't yet have the versatility of the IPv4 address in that they don't have things like geography based or FQDN address but as IPv6 becomes more mainstream this should change.

Subnet Addresses

The Subnet Address type is one that is only used in reference to IPv6 addresses. It represents an IPv6 address subnet. This means that the address will likely be a series of hexadecimal characters followed by a double colon, followed by a "/", and then a number less than 128 to indicate the size of the subnet. An example would be:

fd5e:3c59:35ce:f67e::/64

- The hexadecimal characters represent the IPv6 subnet address.
- The "::" indicates 0's from that point to the left. In an actual address for a computer, the hexadecimal characters that would take the place of these zeros would represent the device address on the subnet.
- /xx, in this case /64 represents the number of bits in the subnet. This will make a range that can potentially include 18,446,744,073,709,551,616 addresses. For those wanting to use English rather than math, that is 18 Quintillion.

Creating a Subnet address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **IPv6 Address**.
4. Input a **Name** for the address object.

5. In the **Type** field, select **Subnet** from the drop down menu.
6. In the **Subnet / IP Range** field, enter the range of addresses in IPv6 format (no spaces)
7. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
8. Input any additional information in the **Comments** field.
9. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	IPv6 Address
Name	IPv6_Guest_user_range
Type	Subnet
Subnet / IP Range	fd5e:3c59:35ce:f67e::/64
Show in Address List	[on]
Comments	

Multicast Addresses

Multicast addressing defines a specific range of address values set aside for them. Therefore all IPv4 multicast addresses should be between 224.0.0.0 and 239.255.255.255.

More information on the concepts behind Multicast addressing can be found in the Multicast Forwarding section.

Multicast IP Range

This type of address will allow multicast broadcasts to a specified range of addresses.

Creating a Multicast IP Range address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category, Multicast Address**
4. Input a **Name** for the address object.
5. Select the **Type, Multicast IP Range** from the dropdown menu.
6. Enter the value for the **Multicast IP Range**
7. Select the **Interface** from the dropdown menu.
8. Enable the **Show in Address List** function

9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: Multicast IP Range Address

The company has a large high tech campus that has monitors in many of its meeting rooms. It is common practice for company wide notifications of importance to be done in a streaming video format with the CEO of the company addressing everyone at once.

The video is High Definition quality so takes up a lot of bandwidth. To minimize the impact on the network the network administrators have set things up to allow the use of multicasting to the monitors for these notifications. Now it has to be set up on the FortiGate firewall to allow the traffic.

- The range being used for the multicast is 239.5.5.10 to 239.5.5.200
 - The interface on this FortiGate firewall will be on port 9
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information

Category	Multicast Address
Name	Meeting_Room_Displays
Type	Multicast IP Range
Multicast IP Range	239.5.5.10-239.5.5.200
Interface	port9
Show in Address List	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall multicast-address
edit "meeting_room_display"
set type multicastrange
set associated-interface "port9"
set start-ip 239.5.5.10
set end-ip 239.5.5.200
set visibility enable
next
end
```

To verify that the address range was added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall multicast-address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Broadcast Subnet

This type of address will allow multicast broadcast to every node on a subnet.

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Multicast Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **Broadcast Subnet** from the drop down menu.
6. In the **Broadcast Subnet** field enter the address and subnet mask according to the format x.x.x.x/x.x.x.x or the short hand format of x.x.x.x/x. (Remember, it needs to be within the appropriate IP range 224.0.0.0 to 239.255.255.255)
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Field	Value
Category	Broadcast Subnet
Name	Corpnet-B
Type	Broadcast Subnet
Broadcast Subnet	224.5.5.0/24
Interface	any
Show in Address List	[on]
Comments	Corporate Network devices - Broadcast Group B

Multicast IP addresses

Multicast uses the Class D address space. The 224.0.0.0 to 239.255.255.255 IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets. The following table lists the reserved multicast address ranges and describes what they are reserved for:

Reserved Multicast address ranges

Reserved Address Range	Use	Notes
224.0.0.0 to 224.0.0.255	Used for network protocols on local networks. For more information, see RFC 1700.	In this range, packets are not forwarded by the router but remain on the local network. They have a Time to Live (TTL) of 1. These addresses are used for communicating routing information.
224.0.1.0 to 238.255.255.255	Global addresses used for multicasting data between organizations and across the Internet. For more information, see RFC 1700.	Some of these addresses are reserved, for example, 224.0.1.1 is used for Network Time Protocol (NTP).
239.0.0.0 to 239.255.255.255	Limited scope addresses used for local groups and organizations. For more information, see RFC 2365.	Routers are configured with filters to prevent multicasts to these addresses from leaving the local system.

Creating multicast security policies requires multicast firewall addresses. You can add multicast firewall addresses by going to **Firewall Objects > Address > Addresses** and selecting **Create New > Multicast Address**. The factory default configuration includes multicast addresses for Bonjour (224.0.0.251-224.0.0.251), EIGRP (224.0.0.10-224.0.0.100), OSPF (224.0.0.5-224.0.0.60), all_hosts (224.0.0.1-224.0.0.1), and all_routers (224.0.0.2-224.0.0.2).

Explicit Proxy Addresses

This category of address is different from the other addresses in that it is not designed to be used in the normal firewall policy configuration. It is intended to be used only with explicit web proxies.

In some respects they can be like a FQDN addresses in that they refer to an alpha-numeric string that is assigned to an IP address, but then goes an additional level of granularity by using additional information and criteria to further specify locations or types of traffic within the website itself. In depth information on Explicit Proxy Addressing can be found in [WAN Optimization](#), but it is worth laying out the steps of how to create an address object for this category.

Creating an Explicit Proxy address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Explicit Proxy Address**.
4. Input a **Name** for the address object.
5. For the **Type** field, select one of the options from the drop down menu.

Within the Explicit Proxy Address category there are 8 types of addresses. Each of these types will have associated field(s) that also need to have values entered to make the object specific to it's address.

Type = URL Pattern

- In the **Host** field, choose from drop down menu
- In the **URL Path Regex** field, enter the appropriate string

Host Regex Match

- In the **Host Regex Pattern** field, enter the appropriate string

URL Category

- In the **Host** field, choose from drop down menu
- In the **URL Category** field, choose from drop down menu

HTTP Method

- In the **Host** field, choose from drop down menu
- In the **Request Method** field, choose from drop down menu

The options are:

- CONNECT
- DELETE
- GET
- HEAD
- OPTIONS
- POST
- PUT
- TRACE

User Agent

- In the **Host** field, choose from drop down menu
- In the **User Agent** field, choose from drop down menu

The options are:

- Apple Safari
- Google Chrome
- Microsoft Internet Explorer or Spartan
- Mozilla Firefox
- Other browsers

HTTP Header

- In the **Host** field, choose from drop down menu
- In the **Header Name** field, enter the appropriate string value
- In the **Header Regex** field, enter the appropriate string value

Advanced (Source)

- In the **Host** field, choose from drop down menu
- In the **Request Method** field, choose from drop down menu (see **HTTP Method** type for option list)
- In the **User Agent** field, choose from drop down menu (see **User Agent** type for option list)
- In the **Header Group** table, create, edit or delete **Header Name** strings and associated **Header Regex** strings

Advance (Destination)

- In the **Host** field, choose from drop down menu
 - In the **Host Regex Pattern** field, enter the appropriate string
 - In the **URL Category** field, choose from drop down menu
6. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
 7. Input any additional information in the **Comments** field.
 8. Press **OK**.

Address Groups

Address groups are designed for ease of use in the administration of the device. If you have a number of addresses or address ranges that will commonly be treated the same or require the same security policies, you can put them into address groups, rather than entering multiple individual addresses in each policy refers to them.

The use of groups is not required. If you have a number of different addresses you could add them individually to a policy and the FortiGate firewall will process them just as quickly and efficiently as if they were in a group, but the chances are that if you have used a group once you could need to use it again and depending on the number of addresses involved entering them individually for each policy can become tedious and the likelihood of an address being missed becomes greater. If you have a number of policies using that combination of addresses it is much easier to add or subtract addresses from the group than to try and remember all of the firewall policies that combination of addresses was used in. With the group, you only have to make the one edit and it is used by any firewall policy using that address group.

Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to Any.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of Any, they can be grouped, even if the addresses involve different networks.

There are 3 Categories of Address groups to choose from:

- IPv4 Group
- IPv6 Group
- Explicit Proxy Group

You cannot mix different categories of addresses within a group, so whether or not it makes sense from an administrative purpose to group certain addresses together, if some are IPv4 and some are IPv6, it cannot be done.

Creating an Address Group

1. Go to **Policy & Objects > Addresses**.
2. Select the down arrow next to **Create New**, select **Address Group**.
3. Choose the **Category**, that is applicable to the proposed selection of addresses.
4. Input a **Group Name** for the address object.

Depending on which **Category** has been chosen the configurations will differ slightly

IPv4 Group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.
3. Select the desired on/off toggle setting for **Static Route Configuration**.

IPv6 Group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.

Explicit Proxy Group

1. Select which Type, either **Source Group** or **Destination Group**.
2. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
3. Select the desired on/off toggle setting for **Show in Address List**.

Irrespective of the Category the groups all have the same final configuration options:

1. Input any additional information in the **Comments** field.
2. Press **OK**.

UUID Support

Syntax:

```
config firewall {address|addres6|addgrp|addgrp6}
  edit 1
    set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
    next
  end
```


Virtual IPs

The mapping of a specific IP address to another specific IP address is usually referred to as Destination NAT. When the Central NAT Table is not being used, FortiOS calls this a Virtual IP Address, sometimes referred to as a VIP. FortiOS uses a DNAT or Virtual IP address to map an External IP address to an IP address. This address does not have to be an individual host, it can also be an address range. This mapping can include all TCP/UDP ports or if Port Forwarding is enabled it will only refer to the specific ports configured. Because, the Central NAT table is disabled by default the term Virtual IP address or VIP will be used predominantly.

Virtual IP addresses are typically used to NAT external or Public IP addresses to internal or Private IP addresses. Using a Virtual IP address between 2 internal Interfaces made up of Private IP addresses is possible but there is rarely a reason to do so as the 2 networks can just use the IP addresses of the networks without the need for any address translation. Using a Virtual IP address for traffic going from the inside to the Internet is even less likely to be a requirement, but it is supported.

Something that needs to be considered when there are multiple Public IP addresses on the external interface(s) is that when a Virtual IP address is used without Port Forwarding enabled there is a reciprocal effect as far as traffic flow is concerned. Normally, on a firewall policy where NAT is enabled, for outgoing traffic the internal address is translated to the Public address that is assigned to the FortiGate, but if there is a Virtual IP address with no port forwarding enabled, then the Internal IP address in the Mapped field would be translated to the IP address configured as the External Address in the VIP settings.

Example

- The assigned External address (WAN1) of the FortiGate unit is 172.12.96.3 with a subnet mask of 255.255.255.128
- There is a Virtual IP address set up to map the external address 172.12.96.127 on WAN1 to the internal IP address of 192.168.1.127
- Port Forwarding is not enabled because you want all allowed traffic going to the external IP address to go to this server.

In this case any outbound traffic from 192.168.1.127 will go out on WAN1 with the IP address of 172.12.96.127 as the source IP address.

In terms of actually using the Virtual IP address, they would be using in the security policies in the same places that other addresses would be used, usually as a Destination Address.

UUID Support for VIP

UUID is now supported in for virtual IPs and virtual IP groups. This includes virtual IPs for IPv4, IPv6, NAT46, and NAT64. To view the UUID for these objects in a FortiGate unit's logs, log-uuid must be set to extended mode, rather than policy-only (which only shows the policy UUID in a traffic log). UUID can only be configured through the CLI

Syntax

```
config sys global
    set log-uuid {disable | policy-only | extended}
end
```



There is another type of address that the term “virtual IP address” commonly refers to which is used in load balancing and other similar configurations. In those cases, a number of devices share a separately created virtual IP address that can be sent to multiple possible devices. In FortiOS these are referred to as Virtual Servers and are configured in the “Load Balance” section.



If Central-NAT is enabled in the CLI the GUI will be different.

Instead of **VIP Type**, the field label will be **DNAT & VIP Type**

Instead of **IPv4** the option will be **IPv4 DNAT**

There will also be the addition setting of **Source Interface Filter**.

Commands to set central-nat:

```
config system settings
  set central-nat [enable | disable]
end
```

Creating a Virtual IP

1. Go to **Policy & Objects > Virtual IPs**.
2. Select **Create New**. A drop down menu is displayed. Select **Virtual IP**.
3. From the **VIP Type** options, choose an applicable type based on the IP addressing involved. Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface.

The available options are:

- **IPv4** - IPv4 on both sides of the FortiGate Unit.
 - **IPv6** - IPv6 on both sides of the FortiGate Unit.
 - **NAT46** - Going from an IPv4 Network to an IPv6 Network.
 - **NAT64** - Going from an IPv6 Network to an IPv4 Network.
4. In the **Name** field, input a unique identifier for the Virtual IP.
 5. Input any additional information in the **Comments** field.

In the **Network** section

6. If an IPv4 type of Virtual IP, select the **Interface** setting.

Using the dropdown menu for the Interface Field, choose the incoming interface for the traffic.

The IPv4 VIP Type is the only one that uses this field. This is a legacy function from previous versions so that they can be upgraded without complicated reconfiguration. The External IP address, which is a required field, tells the unit which interface to use so it is perfectly acceptable to choose **"any"** as the interface. In some configurations, if the Interface field is not set to **"any"** the Virtual IP object will not one of the displayed options when choosing a destination address.

7. Configure the **Source Interface Filter** (if available)

If needed, toggle the setting on. This will cause the field with a "+" symbol in it to appear. Once the field is selected, a single or multiple interfaces can be selected from the window that slides out from the right.

8. Configure the **External IP Address/Range**.

There are two fields. If there is a single IP address, use that address in both fields. This will be the address on the outside of the network that is usually the public address of the server. The format of the address will depend on the **VIP Type** option that was selected.

9. Configure the **Mapped IP Address/Range**. This will be the address that the traffic is being directed to.

There are two fields. If there is a single IP address, use that address in both fields. The format of the address will depend on the **VIP Type** option that was selected.

10. Disable/Enable the **Source Address Filter**.

If only specific IP addresses are allowed to be the source address for traffic using the VIP, enable the **Source Address Filter**. To add an allowed address select **Create New**. The value for the address field for the Source Address Filter can be formatted in three different ways.

- **Source IP** - Use the standard format for a single IP address based on whether it's IPv4 or IPv6
- **Range** - Enter the first and last members of the range
- **Subnet** - Enter the IP address of the broadcast address for the subnet.

11. Disable/Enable **Port Forwarding**. If only the traffic for a specific port or port range is being forwarded, enable this setting.

i. Select the **Protocol**.

Depending on which Virtual IP type is being configured there can be one of up to 4 different protocols being forwarded.

- **IPv4** can forward: TCP, UDP, SCTP or ICMP
- **IPv6** can forward: TCP, UDP, or SCTP
- **NAT46** can forward: TCP or UDP
- **NAT64** can forward: TCP or UDP

ii. Configure the **External Service Port**.

This will be the listening port that the traffic is being sent to. If **ICMP** was selected, there will not be any port options available. This is because only one internal address will be able to respond to ICMP requests. For the other options there will be 2 field to configure. The start and the end of the port range. If only a single port is being configured, enter the same value in both fields.

iii. Configure the setting **Map to Port**.

This will be the listening port on the device on the internal side of the network. It does not have to be the same as the **External Service Port**. There will be 2 field to configure. The start and the end of the port range. If only a single port is being configured, enter the same value in both fields.

12. Press **OK**.

Example

This example is for a VIP that is being used to direct traffic from the external IP address to a webserver on the internal network. The webserver is for company use only. The company's public facing webserver already used port 80 and there is only one IP external IP address so the traffic for this server is being listened for on port 8080 of the external interface and being sent to port 80 on the internal host.

Field	Value
VIP Type	IPv4
Name	Internal_Webserver
Comments	Webserver with Colaboration tools for Corporate employees
Interface	Any
External IP Address/Range	172.13.100.27 <this would normally be a public IP address>
Mapped IP Address/Range	192.168.34.150
Source Address Filter	<list of IP addresses of remote users>
Port Forwarding	enabled
Protocol	TCP
External Service Port	8080 - 8080
Map to Port	80 - 80

Dynamic VIP according to DNS translation

When a dynamic virtual IP is used in a policy, the dynamic DNS translation table is installed along with the dynamic NAT translation table into the kernel. All matched DNS responses will be translated and recorded regardless if they hit the policy. When a client request hits the policy, dynamic NAT translation will occur if it matches a record, otherwise the traffic will be blocked.

Syntax

```
config firewall vip
  edit "1"
    set type dns-translation
    set extip 192.168.0.1-192.168.0.100
    set extintf "dmz"
    set dns-mapping-ttl 604800
    set mappedip "3.3.3.0/24" "4.0.0.0/24"
  end
end
```

Virtual IP Groups

Just like other address, Virtual IP addresses can be organized into groups for ease of administration. If you have multiple virtual IPs that are likely to be associated to common firewall policies rather than add them individually to each of the policies you can add the instead. That way, if the members of the group change then any changes made to the group will propagate to all of the policies using that group.

When using a Virtual IP address group the firewall policy will take into account all of the configured parameters of the Virtual IPs: IP addresses, Ports and port types.

Creating a Virtual IP Group

1. Go to **Policy & Objects > Virtual IPs**.
2. Select **Create New**. A drop down menu is displayed. Select **Virtual IP**.
3. Select the **Type** fo VIP group you wish to create.
The options available are:
 - **IPv4** - IPv4 on both sides of the FortiGate Unit.
 - **IPv6** - IPv6 on both sides of the FortiGate Unit.
 - **NAT46** - Going from an IPv4 Network to an IPv6 Network.
 - **NAT64** - Going from an IPv6 Network to an IPv4 Network.

Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface. The options will be:

4. Enter a unique identifier for the group in the **Name** field.
5. Enter any additional information in the **Comments** field.
6. Use the drop-down menu of the **Interface** field to select the interface if all of the VIPs are on the same interface. If any of the VIPS are on different interfaces or if any of them are associated with the "any" option, choose the any option for the group.
7. Select anywhere in the **Members** field to bring forth the pane of potential members for selection to the group.
8. Press **OK**.

Configuring IP pools

A IP pool is essentially one in which the IP address that is assigned to the sending computer is not known until the session is created, therefore at the very least it will have to be a pool of at least 2 potential addresses. A quick example would be an IP pool for users of a VPN. IP pools are based upon the version of IP determined by the interface that they are associated with so as expected there are two types of IP pools that can be configured:

- IPv4 Pool
- IPv6 Pool

Because of the differences in the configuration for the two types of pools, instructions for configuring them will be done separately.

Creating a IPv4 Pool

1. Go to **Policy & Objects > IP Pools**.
2. Select **Create New**.
3. In the IP Pool Type field choose **IPv4 Pool**
4. Enter a name in the **Name** field for the new service
5. Include any description you would like in the **Comments** field
6. In the **Type** field choose between:
 - **Overload**
 - **One-to-One**
 - **Fixed Port Range**
 - **Port Block Allocation**

At this point the configurations can start to differ based on the type of type of pool.

For more information on the different types of IP pools, check [IP Pools](#) in the Concepts section.

Overload

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Enable the **ARP Reply** field by making sure there is a check in the box
9. Select **OK**

Overload Example for GUI

In this example, the Sales team needs to connect to an Application Service Provider that does the accounting for the company. As a security measure, the ASP only accepts traffic from a white list of IP addresses. There is 1 public IP address of the company on that list. The Sales team consists of 40 people, so they need to share. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	Sales_Team
Comments	For the Sales team to use to connect to the Accounting ASP
Type	Overload (This is the default)
External IP Range	10.23.56.20 - 10.23.56.20
ARP Reply	enabled

Overload Example for CLI

```
config firewall ippool
edit Sales_Team
set comments "For the Sales team to use to connect to the Accounting ASP"
```

```
set type overload
set startip 10.23.56.20
set endip 10.23.56.20
set arp-reply enable
set arp-intf wan1
end
```

One-to-one

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Enable the **ARP Reply** field by making sure there is a check in the box.
9. Select **OK**

One-to-one Example for GUI

In this example, the external IP address of the mail server is part of a range assigned to the company but not the one that is assigned to the Internet facing interface. A VIP has been set up but in order to properly resolve Reverse DNS lookups the mail server always has to use a specific IP address. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	Mail-Server
Comments	So the the correct IP address is resolved on Reverse DNS look ups of the mail server.
Type	One-to-one
External IP Range	10.23.56.21 - 10.23.56.21
ARP Reply	enabled

One-to-one Example for CLI

```
config firewall ippool
edit Mail-Server
set comments "So the the correct IP address is resolved on reverse DNS look ups of
the mail server."
set type one-to-one
set startip 10.23.56.21
set endip 10.23.56.21
set arp-reply enable
set arp-intf wan1
end
```

Fixed Port Range

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Fort the **Internal IP Range** fields, enter the lowest and highest addresses in the range.

9. Enable the **ARP Reply** field by making sure there is a check in the box
10. Select **OK**

Fixed Port Range Example for GUI

In this example, the company has a range of 10 IP address that they want to be used by employees on a specific subnet for NATing. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	IPPool-3
Comments	IP range to be used by outgoing traffic
Type	Fixed Port Range
External IP Range	10.23.56.22 - 10.23.56.31
Internal IP Range	192.168.23.1 - 192.168.23.254
ARP Reply	enabled

Fixed Port Range Example for CLI

```
config firewall ippool
edit IPPool-3
    set comments "So the the correct IP address is resolved on reverse DNS look ups of
    the mail server."
    set type fixed-port-range
    set startip 10.23.56.22
    set endip 10.23.56.31
    set source-startip 192.168.23.1
    set source-endip 192.168.23.254
    set arp-reply enable
    set arp-intf wan1
end
```

Port Block Allocation

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. In the **Block Size** field, either type in the value or use the up or down arrows to set the value of the block size.
9. In the **Blocks Per User** field, either type in the value or use the up or down arrows to set the value for the number of blocks per user.
10. Enable the **ARP Reply** field by making sure there is a check in the box
11. Select **OK**

Port Block Allocation Example for GUI

In this example, a small ISP is setting up NATing for its clients, but to be fair it is putting some restrictions on the number of connections each client can have so that no one hogs all of the possible ports and addresses. The external interface is port12.

Field	Value
IP Pool Type	IPv4 Pool
Name	Client-IPPool
Comments	IP Pool for clients to access the Internet
Type	Port Block Allocation
External IP Range	10.23.75.5 - 10.23.75.200
Block Size	64
Blocks Per User	8
ARP Reply	enabled

Port Block Allocation Example for CLI

```
config firewall ippool
edit Client-IPPool
    set comments "IP Pool for clients to access the Internet"
    set type port-block-allocation
    set startip 10.23.75.5
    set endip 10.23.75.200
    set block-size 64
    set num-blocks-per-user 8
    set permit-any-host disable
    set arp-intf wan1
    set arp-reply enable
    set arp-intf port12
end
```

Creating a IPv6 Pool

1. Go to **Policy & Objects > IP Pools**.
2. Select **Create New**.
3. In the IP Pool Type field choose **IPv6 Pool**
4. Enter a name in the **Name** field for the new service
5. Include any description you would like in the **Comments** field
6. For the **External IP Range** fields, enter the lowest and highest addresses in the range.

IPv6 Example for GUI

In this example, there is a similar situation to the One-to-one example earlier. There is a mail server that needs to be resolved to a specific IP address in Reverse DNS look-ups. The difference in this case is the company is an early adopter of IPv6 connectivity to the Internet.

Field	Value
IP Pool Type	IPv6 Pool
Name	Mail-svr-ipv6
Comments	Registered IPv6 address for mail server
External IP Range	fd2f:50ec:cdea:0663::1025 - fd2f:50ec:cdea:0663::1025

Port Block Allocation Example for CLI

```
config firewall ippool6
edit Mail-svr-ipv6
set comments "Registered IPv6 address for mail server"
set startip fd2f:50ec:cdea:663::102
set endip fd2f:50ec:cdea:663::1025
end
```

Services

While there are a number of services already configured within FortiOS, the firmware allows for administrators to configure their own. The reasons for doing this usually fall into one or more of the following categories:

- The service is not common enough to have a standard configuration
- The service is not established enough to have a standard configuration
- The service has a standard port number but there is a reason to use a different one:
 - Port is already in use by another service
 - For security reasons, want to avoid standard port

When looking at the list of preconfigured services it may seem like there are a lot, but keep in mind that the theoretical limit for port numbers is 65,535. This gives a fairly good sized range when you are choosing what port to assign a service but there are a few points to keep in mind.

- Most of the well known ports are in the range 0 - 1023
- Most ports assigned by the Internet Corporation for Assigned Names and Numbers (ICANN) will be in the 1024 - 49151 range
- Port numbers between 49,152 and 65,535 are often used for dynamic, private or ephemeral ports.

There are 3 Service objects that can be added and configured:

- Categories
- Services
- Service Groups

Categories

In order to make sorting through the services easier, there is a field to categorize the services. Because selecting a category is part of the process for creating a new service, the configuration of categories will be explained first.

The services can be sorted into the following groups:

- General
- Web Access
- File Access
- Email
- Network Services
- Authentication
- Remote Access
- Tunneling
- VoIP, Messaging and Other Applications
- Web Proxy
- Uncategorized

The categories are for organization purposes so there is not many settings when creating a new one.

Creating a new Service Category

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Category**
3. Input a **Name** for the category..
4. Input any additional information in the **Comments** field.
5. Press **OK**.

Example

You plan on adding a number of devices such as web cameras that will allow the monitoring of the physical security of your datacenter. A number of non-standard services will have to be created and you would like to keep them grouped together under the heading of "Surveillance"

Example of a New Category in the GUI

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Category**.
2. Fill out the fields with the following information

Field	Value
Name	Surveillance
Comments	For DataCenter Surveillance Devices

3. Select **OK**.

Example of a New Category in the CLI

Enter the following CLI command:

```
config firewall service category
edit Surveillance
set comment "For DataCenter Surveillance Devices"
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Select the Category Settings icon . A listing of the categories should be displayed.
2. Enter the following CLI command:

```
config firewall service category
show
```

This should bring up all of the categories. Check to see that the new one is displayed.

Configuring a new service

Occasionally, the preconfigured list of services will not contain the needed service. There are a few variations in the creation of a service depending upon the protocol type, but the first steps in the creation of the service are common to all the variations.

To create a new service:

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Service**
3. Enter a name in the **Name** field for the new service
4. Include any description you would like in the **Comments** field
5. In the **Service Type** field choose between **Firewall** and **Explicit Proxy**. For the purposes of this chapter **Firewall** will always be chosen. **Explicit Proxy** services covered are in the WAN-OPT Handbook .
6. Enable the toggle in the **Show in Service List**. If you can't see the service when you need to select it, it serves very little purpose.
7. For the **Category** field, choose the appropriate category from the **Category** drop down menu. If none is chosen, the **Uncategorized** option will be chosen by default.

Protocol Options

This is the section where the configuration options of the service will differ depending on the type of protocol chosen. (The Step numbers will all continue on from the common step sequence)

TCP/UDP/SCTP

8. For the **Protocol Type** field, choose **TCP/UDP/SCTP** from the drop down menu
9. In the **IP/FQDN** field, an IP address or Fully Qualified Domain name can be entered if there is to be a specific destination for the service
10. Configure the **Destination Port** by:

- Select from the drop down menu, **TCP**, **UDP** or **SCTP**
 - Enter the low end to the port range in the field indicated by grayed out **Low**.
 - Enter the high end of the port range in the field indicated by grayed out **High**. If there is only a single port in the range **High** can be left empty
 - Multiple ports or port ranges can be added by using the "+" at the beginning of the row
 - Rows can be removed by using the trash can symbol at the end of the row
11. If required, you can **Specify Source Ports** for the service by enabling the toggle switch.
- The **Src Port** will match up with a **Destination Port**
 - **Src Ports** cannot be configured without there being a value for the **Destination Port**
 - The same rules for configuring the **Destination Ports** applies to the **Src Ports**
12. Select **OK** to confirm the configuration

Example

Example settings for a TCP protocol service. In this case, it is for an administrative connection to web servers on the DMZ. The protocol used is HTTPS which would normally use port 443, but that is already in use by another service such as Admin access to the firewall or an SSL-VPN connection.

Field	Value
Name	Example.com_WebAdmin
Comments	Admin connection to Example.com Website
Service Type	Firewall
Show in Service List	enabled
Category	Web Access
Protocol Options	
Protocol Type	TCP/UDP/SCTP
IP/FQDN	<left blank>
Destination Port	<ul style="list-style-type: none"> • Protocol: TCP • Low: 4300 • High: <left blank>
Specify Source Ports	<disabled>

Creating a new TCP/UDP/SCTP service in the CLI

The following is the creation of the same service using the command line.

```
config firewall service custom
edit Example.com_WebAdmin
set comment "Admin connection to Example.com Website"
```

```

        set category Web Access
        set protocol TCP/UDP/SCTP
        set tcp-portrange 4300
    end
end

```

ICMP / ICMP6

8. For the **Protocol Type** field, choose **ICMP** or **ICMP6** from the drop down menu
9. In the **Type** field enter the appropriate type number based on the information found in ["ICMP Types and Codes" on page 1](#) or in ["ICMPv6 Types and Codes" on page 63](#), depending on whether the **Protocol Type** is **ICMP** or **ICMPv6**
10. In the **Code** field enter the appropriate code number for the type, if applicable, based on the information found in ["ICMP Types and Codes" on page 1](#) or in ["ICMPv6 Types and Codes" on page 63](#), depending on whether the **Protocol Type** is **ICMP** or **ICMPv6**
11. Select **OK** to confirm the configuration

Example

Example settings for an ICMP.service. In this case it has been set up for some special testing of ICMP packets.

Field	Value
Name	ICMP test #4
Comments	For testing of proprietary network scanner
Service Type	Firewall
Show in Service List	enabled
Category	Network Services
Protocol Options	
Protocol Type	ICMP
Type	7
Code	<left blank>

Creating a new ICMP service in the CLI

The following is the creation of the same service using the command line.

```

config firewall service custom
    edit ICMP test4
        set comment "For testing of proprietary network scanner"
        set category Network Services
        set protocol ICMP
        set icmptype 7
    end

```

end

IP

8. For the **Protocol Type** field, choose **IP** from the drop down menu
9. In the **Protocol Number** field enter the numeric value based on the information found in ["Protocol Number" on page 74](#)
10. Select **OK** to confirm the configuration

Example

Example settings for an IP.service. In this case it has been set up to communicate via an old protocol called QNX

Field	Value
Name	QNX
Comments	For QNX communications to the Development Lab
Service Type	Firewall
Show in Service List	enabled
Category	Uncategorized
Protocol Options	
Protocol Type	IP
Protocol Number	106

Creating a new ICMP service in the CLI

The following is the creation of the same service using the command line.

```
config firewall service custom
edit ICMP test4
    set comment "For QNX communications to the Development Lab "
    set protocol IP
    set icmptype 106
end
end
```



In the CLI examples, the fields for **Show in Service List**, **Service Type** and in the example for IP, **Category** were not set because the values that they would have been set to were the default values and were already correctly set.

Example Scenario: Using FortiGate services to support Audio/Visual Conferencing



The feature, and the transmitting of data for the purpose of, Tele-conferencing or Audio/Visual Conferencing is covered by a number of standards:

- The IETF standard known as the Binary Floor Control Protocol (BFCP).
- RFC 4582, for SIP-based video devices
- The ITU standard H.239 (for H.323-based video devices)

While these standards have been set up by various authoritative bodies and can take place on different layers of the OSI model, they share common requirements that are addressed by the FortiGate firewall's ability to manage the traffic and the protocols involved. This means that the same ability that make the device RFC 4582 compliant makes it compliant with H.239 as well.

To demonstrate how services and service groups are used we show the setup of a firewall that will need to support the connectivity of a video conferencing unit. The FortiGate does not manipulate or change the content of the traffic but it does allow for the traffic to pass through the device. In this case it allow for only the needed traffic to pass through the device so as to allow the functionality of Audio Visual Conference call but not to allow other traffic through.

The theoretical location for this scenario is a hospital that hosts conferences and lectures from doctors from all over the world, sometimes from multiple locations, using video conferencing technology such as a Polycom Video Conference system. There is a special room set up with dedicated Ethernet connectivity to the Internet. A hospital has a lot of sensitive information going over its network so the setup has to be secure to prevent any chance of penetration.

The approach is fairly simple. The conference room has a dedicated port on the FortiGate (port #7) and its own LAN. We will assume that the interface has already been configured properly. Video conference traffic can come from the Internet to the Polycom in that room and traffic can get out to the Internet, but traffic going to other areas of the hospital network have to go through the FortiGate and traffic going from the Video Conference LAN is thoroughly filtered.

To give an idea of how extensive this can be, we will use an extreme case and include just about all of the services that could be commonly used in one of these setups. The protocols listed here may differ from other setups. It will depend on which features are being used and which equipment is within the network. Always check the documentation that comes with the set up before opening ports into your network.

VIP

In this particular case there is an IP address set aside for the conferencing system so a separate VIP is not needed for every port. One Virtual IP will be created for the system and then only the approved of protocols will be allows through the firewall.

Name	Vid-Conf_Room216
External Interface	wan1

External IP Address/Range	256.87.212.51 – 256.87.212.51
Mapped IP Address/Range	192.168.7.25 – 192.168.7.25
Port Forwarding	not selected

Creating an address for the subnet

In the same way that the VIP was created to identify and direct incoming traffic an address should be created to identify the addresses of computer that will be in the Conference room. This included computers on the LAN as well as the Teleconferencing equipment.

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Category	Address
Name	Port7_subnet
Type	Subnet
Subnet/IP Range	192.168.7.0/255.255.255.0
Interface	port7
Show in address list	checked

Configuring the services

Services already created:

The following are standard services that have already been created by default:

HTTP	TCP 80
SNMP	TCP 161-162/UDP 161-162
LDAP	TCP 389
HTTPS	TCP 443
SYSLOG	UDP 514

Existing Services to be edited:

There are a few services that have already been created for you, but they need to be expanded to accommodate the list of protocols listed for this scenario.

The default h323 contains:

- TCP 1503
- UDP 1719
- TCP 1720

We need to add:

- TCP1719

The default SIP contains:

- UDP 5060

We need to add:

- TCP 5060

H323 service

1. Go to **Policy & Objects > Objects > Services**.
2. Scroll down to the section: **VoIP, Messaging & Other Applications**.
3. Select **H323**.
4. Select **Edit**.
5. In the Protocol section add the additional protocol:

Protocol Type	TCP
Destination port /Low	1719

6. Select **OK** to save.

SIP service

1. Go to **Policy & Objects > Objects > Services**.
2. Scroll down to the section: **VoIP, Messaging & Other Applications**.
3. Select **SIP**.
4. Select **Edit**.
5. In the Protocol section add the additional protocol:

Protocol Type	TCP
Destination port /Low	5060

6. Select **OK** to save.

Custom Services that need to be created

There are a number of possible services that may need to be added from scratch rather than editing existing ones. While it is possible to create a single custom service that contains all of the open ports needed, it makes more sense to make this modular in case only a small subset of the service needs to be added to another policy.

Polycom API

1. Go to **Policy & Objects > Objects > Services**.
2. Select **Create New**.
3. Fill in the fields of the new service with the following information:

Name	Polycom API
Service Type	Firewall
Category	VoIP, Messaging & Other
Protocol Type	TCP/UDP/SCTP
Protocol	TCP/UDP/SCTP
Protocol	TCP
Destination Port - Low:	24
Destination Port - High:	<leave blank>

4. Select **OK**.

Polycom Endpoints

1. Go to **Policy & Objects > Objects > Services**.
2. Select **Create New**.
3. Fill in the fields of the new service with the following information:

Name	Polycom Endpoints
Service Type	Firewall
Category	VoIP, Messaging & Other
Protocol Type	TCP/UDP/SCTP
Protocol	TCP
Destination - Low:	3230
Destination - High:	3253

4. Select **OK**.

Other Services to add in the same way:

Name of Service	Category	Protocol & Port #
LDAP secure communications	Authentication	TCP 636
Win 2000 ILS Registration	Network Services	TCP 1002
Gatekeeper discovery	VoIP, Messaging & Other Applications	TCP 1718
Audio Call Control	VoIP, Messaging & Other Applications	TCP 1731
Polycom proprietary Global directory data	VoIP, Messaging & Other Applications	TCP 3601
Polycom People+Content	VoIP, Messaging & Other Applications	TCP 5001
HTTP Server Push	Web Access	

Creating the Service Group

1. Go to **Firewall Objects > Service > Groups**.
2. Select **Create New**.
3. Build the Service group by filling in the fields with the following information

Group Name	A-V_Conference
Type	Firewall
Members (click in the drop down menu to add the following services)	<ul style="list-style-type: none"> • HTTP • SNMP • LDAP • HTTPS • SYSLOG • Polycom API • Polycom Endpoints • LDAP secure communications • Win 2000 ILS Registration • Gatekeeper discovery • Audio Call Control • Polycom proprietary Global directory data • Polycom People+Content • HTTP Server Push

Creating the IPS Security Profile

This is by no means the only way to set up this IPS filter, but it is the way that the fictional System Administrator wants it set up. Yours may be different.

1. Go to Security **Profiles > Intrusion Protection > IPS Sensors**.
2. Create a new sensor.

Name	A-V_Conference-incoming
-------------	-------------------------

3. Select **OK**.
4. In the newly created sensor, create a new IPS filter.

Sensor Type	Filter Based
Filter Options	Advanced
Severity	<ul style="list-style-type: none"> • Critical • High • Medium • Low
Target	Server
OS	Windows
Application	<ul style="list-style-type: none"> • IIS • other
Protocol Use the [Show more...] option	<ul style="list-style-type: none"> • HTTP • LDAP • SIP • SSL • H323
Packet logging	enabled

Based on these filters there should be somewhere in the neighborhood of 750 signatures that the FortiGate will run traffic against in the IPS engine.

Policies

Incoming Policy

A policy has to be made to allow the traffic to come in from the Internet to connect to the Tele-conferencing server equipment.

1. Go to **Policy & Objects > Policy > IPv4**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Policy Type	Firewall
Policy Subtype	Address
Incoming Interface	wan1
Source Address	all

Outgoing Interface	port7
Destination Address	Vid-Conf_Room216
Schedule	always
Service	A-V_Conference
Action	ACCEPT
Enable NAT	<not enabled>
Logging Options	Logging is a good idea but how much will depend on storage capabilities.
Security Profiles	Turn on IPS and choose "A-V_Conference-incoming"
Traffic Shaping, Web cache, WAN Optimization, Disclaimer:	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

4. Select **OK**.

The policy will then need to be put in the correct position in the sequence of the policies. Because it is a rather focused policy it should be acceptable to place it near the top of the policy order sequence.

Outgoing Policy

A policy has to be made to allow the traffic to leave from the subnet in the conference room to the Internet, not only for the traffic for the Tele-conferencing equipment but for normal traffic of users on the Internet such as web research and email. The traffic is outgoing so there is less of a need for an Intrusion Protection System filter, but check with the network architect in case there is a need for using one of the other security profiles.

1. Go to **Policy & Objects > Policy > IPv4**.
2. Select **Create New**.
3. Fill out the fields with the following information:

Policy Type	Firewall
Policy Subtype	Address
Incoming Interface	port7
Source Address	Port7_subnet
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	any
Action	ACCEPT

Enable NAT	enabled Use Destination Interface Address
Logging Options	Logging is a good idea but how much will depend on storage capabilities.
Security Profiles	<see above>
Traffic Shaping, Web cache, WAN Optimization, Disclaimer:	The use of these features will depend on your network environment and should be decided by the network architect, as the decision will largely be based on network bandwidth, usage and importance of Video conferencing compared to other traffic.

4. Select **OK**.

The policy will then need to be put in the correct position in the sequence of the policies.

Specific Addresses in TCP/UDP/SCTP

In the TCP/UDP/SCTP services it is also possible to set the parameter for a specific IP or Fully Qualified Domain Name address. The IP/FQDN field refers to the destination address of the traffic, not the source. This means for example, that you can set up a custom service that will describe in a policy the TCP traffic over port 80 going to the web site example.com, but you cannot set up a service that describes the TCP traffic over port 80 that is coming from the computer with the address 192.168.29.59.

Service Groups

Just like some of the other firewall components, services can also be bundled into groups for ease of administration.

Creating a ServiceGroup

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Service Group**
3. Input a **Group Name** to describe the services being grouped
4. Input any additional information in the **Comments** field.
5. Choose a **Type** of group. The options are **Firewall** or **Explicit Proxy**.
6. Add to the list of **Members** from the drop down menu. Using the + sign beside the field will allow the addition of multiple services.
7. Press **OK**.

Example

Example of a New Service Group:

Field	Value
Group Name	Authentication Services
Comments	Services used in Authentication

Field	Value
Type	Firewall
Members	<ul style="list-style-type: none">• Kerberos• LDAP• LDAP_UDP• RADIUS

Firewall schedules

Firewall schedules control when policies are in effect. When you add a security policy on a FortiGate unit you need to set a schedule to determine the time frame in which that the policy will be functioning. While it is not set by default, the normal schedule would be always. This would mean that the policy that has been created is always function and always policing the traffic going through the FortiGate. The time component of the schedule is based on a 24 hour clock notation or military time as some people would say.

There are two types of schedules: One-time schedules and recurring schedules.

One-Time schedules are in effect only once for the period of time specified in the schedule. This can be useful for testing to limit how long a policy will be in effect in case it is not removed, or it can be used for isolated events such as a conference where you will only need a temporary infrastructure change for a few days.

The time frame for a One-time schedule is configured by using a start time which includes, Year | Month | Day | Hour | Minute and a Stop time which includes the same variables. So while the frequency of the schedule is only once it can last anywhere from 1 minute to multiple years.

Recurring schedules are in effect repeatedly at specified times of specified days of the week. The Recurring schedule is based on a repeating cycle of the days of the week as opposed to every x days or days of the month. This means that you can configure the schedule to be in effect on Tuesday, Thursday, and Saturday but not every 2 days or on odd numbered days of the month.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next.

Creating a recurring schedule object

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule**.
3. From the **Type** options, choose **Recurring**.
4. Input a **Name** for the schedule object.
5. From the **Days** options, choose the day of the week that you would like this schedule to apply to. The schedule will be in effect on the days of the week that have a check mark in the checkbox to the left of the name of the weekday.
6. Choose a **Start Time**.
The **Start Time** is composed of two fields, **Hour** and **Minute**. Think of setting the time for a digital clock in 24 hour mode. The **Hour** value can be an integer from 0 and 23. The **Minute** value can be from 0 to 59. 0 and 0 would be midnight at the start of the day and 23 and 59 would be one minute to midnight at the end of the day. The value can be entered by keyboard or by using the up and down arrows in the field to select the value.
7. Choose a **Stop Time**.
Configuration is the same as **Start Time**.
8. Press **OK**.



Because recurring schedules do not work with DENY policies, the strategy when designing a schedule should *not* be to determine when users cannot access a policy but to build the schedules around when it *is* possible to access the policy.

Creating a One-time schedule object

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule**.
3. From the **Type** options, choose **One-time**.
4. Input a **Name** for the schedule object.
5. Choose a **Start Date**.
Selecting the field with the mouse will bring up a interactive calendar graphic that will allow the user to select the date. The date can also be typed in using the format YYYY/MM/DD.
6. Choose an End Date.
Configuration is the same as **Start Date**.
7. Choose a **Start Time**.
The **Start Time** is composed of two fields, **Hour** and **Minute**. Think of setting the time for a digital clock in 24 hour mode. The **Hour** value can be an integer from 0 and 23. The **Minute** value can be from 0 to 59. 0 and 0 would be midnight at the start of the day and 23 and 59 would be one minute to midnight at the end of the day. The value can be entered by keyboard or by using the up and down arrows in the field to select the value.
8. Choose a **Stop Time**.
Configuration is the same as **Start Time**.
9. Enable/Disable **Pre-expiration event log**.
This configures the system to create an event log 1 to 100 days before the **End Date** as a warning in case the schedule needs to be extended.
10. If the **Pre-expiration event log** is enabled, set the value for **Number of days before**.
11. Press **OK**.

Example

You want to schedule the use of Skype to only between noon (12:00) and 1 p.m. (13:00).

You could create a schedule that allows Skype traffic:

- Starting at Hour:12 and Minute: 00
- Stopping at Hour:13 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Or you could have a schedule that blocks Skype traffic:

- Starting at Hour:13 and Minute: 00 (and goes to the next day)
- Stopping at Hour:12 and Minute: 00
- Set for days of the week: Sunday | Monday |Tuesday |Wednesday | Thursday | Friday | Saturday

Either way is effective for the task but other factors may make one method work better than another in certain situations or it could be just a preference in approach.

Schedule expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this option is set to disable.

A few further settings are needed to make this work.

```
config firewall policy
  edit ID
    set firewall-session-dirty check-new
  end

config system settings
  set firewall-session-dirty check-policy-option
end
```

Firewall-session-dirty setting

The firewall-session-dirty setting has three options

check-all	CPU flushes all current sessions and re-evaluates them. [default]
check-new	CPU keeps existing sessions and applies policy changes to new sessions only. This reduces CPU load and the possibility of packet loss.
check-policy-option	Use the option selected in the firewall-session-dirty field of the firewall policy (check-all or check-new, as above, but per policy).

Schedule Groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. The schedule parameter in the policy configuration does not allow for the entering of multiple schedules into a single policy so if you have a combination of time frames that you want to schedule the policy for then the best approach, rather than making multiple policies is to use a schedule group.

Creating a recurring schedule object

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule Group**
3. Input a **Name** for the schedule object.
4. In the **Members** field, select the "+" to bring forth the panel for selecting entries.
5. Press **OK**.

Example

Your Internet policy allows employees to visit Social Media sites from company computers but not during what is considered working hours. The offices are open a few hours before working hours and the doors are not locked until a few hours after official closing so work hours are from 9 to 5 with a lunch break from Noon to 1:00 p.m.

Your approach is to block the traffic between 9 and noon and between 1:00 p.m. and 5:00 p.m. This means you will need two schedules for a single policy and the schedule group handles this for you. Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule Expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this is set to `disable`.



FORTINET®

High Performance Network Security



Copyright© 2018 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.